

TR-95-09

ACOUSTO-OPTICAL FREQUENCY SHIFTERS IN DUAL-CORE
OPTICAL FIBRES

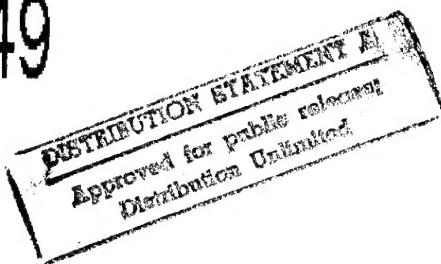
AFOSR 91-0387

Final Report
March 1995

T. A. Birks, P. St.J. Russell, D. O. Culverhouse, S. G. Farwell & C. N. Pannell

Optoelectronics Research Centre
University of Southampton
Hampshire SO9 5NH
United Kingdom

19950427 149



DTIC QUALITY CERTIFIED 5

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

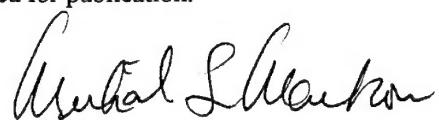
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Washington Headquarters Services Directorate for Information Operations and Analysis, 1215 Jefferson Davis Highway, Suite 1200, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20585.

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED	
	March 1995	Final Report 30 Sep 93-29 Mar 95	
4. TITLE AND SUBTITLE		5. FUNDING NUMBER	
Acousto-Optical Frequency Shifters in Dual-Core Optical Fibres		G/AFOSR-91-0387 PR/0900/10 PR/4200/90	
6. AUTHOR(S)		7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	
T.A.Birks, P.St.J.Russell, D.O.Culverhouse S.G. Farwell and C.N. Pannell.		Optoelectronics Research Centre University of Southampton Hants, SO9 5NH, United Kingdom	
8. PERFORMING ORGANIZATION REPORT NUMBER		9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	
		EOARD 223/231 Old Marylebone Road London, NW1 5TH, United Kingdom	
10. SPONSORING/MONITORING AGENCY REPORT NUMBER		11. SUPPLEMENTARY NOTES	
TL-95-09			
12a. DISTRIBUTION/AVAILABILITY STATEMENT		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words)		14. SUBJECT TERMS	
The aim of this project is the design and implementation of a practical acousto-optic frequency shifter compatible with single-mode optical fibre. The frequency shifter was originally to be based on a special fibre with two dissimilar cores. Normally, little light couples between the cores, but in the presence of an appropriate flexural acoustic wave a large transfer of optical power, with a frequency shift, takes place. This interaction has been analysed theoretically, a dual-core fibre frequency shifter designed, and its component parts assembled.		15. NUMBER OF PAGES 73	
Progress towards this initial experiment was hampered by fibre fabrication problems. In the meantime, an alternative type of frequency shifter was pursued. Here the acousto-optic interaction takes place between two modes in the narrow uniform waist of a special "null" fused taper coupler. This has so many compelling advantages that the earlier dual-core work was put to one side. A prototype null coupler device was built and tested, and immediately satisfied some of the project's goals. The design promises to be able to satisfy the remainder also.		16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT		18. SECURITY CLASSIFICATION OF THIS PAGE	
19. SECURITY CLASSIFICATION OF ABSTRACT		20. LIMITATION OF ABSTRACT	

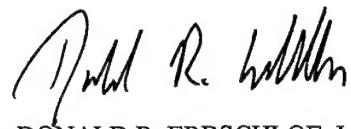
TR-95-09

This report has been reviewed and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

This technical report has been reviewed and is approved for publication.



MICHAEL S. MARKOW, Lt Col, USAF
Chief, Aerospace Electronics



DONALD R. ERBSCHLOE, Lt Col, USAF
Chief, International Programs

ACOUSTO-OPTICAL FREQUENCY SHIFTERS IN DUAL-CORE OPTICAL FIBRES

AFOSR 91-0387

Final Report
March 1995

Contents	page
1. SUMMARY	4
2. INTRODUCTION	4
2.1 Project goals	4
2.2 Realisation of these goals	6
3. ANALYSIS OF THE ACOUSTO-OPTIC INTERACTION IN FIBRES	8
3.1 Description	8
3.2 Passive behaviour	9
3.3 Acousto-optic interaction	10
3.4 Frequency shifter output powers	11
3.5 Maximum conversion	13
3.6 Frequency shift	14
3.7 Bandwidth	15
3.8 Compensation for one parameter by another	16
3.9 Carrier and image sideband suppression	16
3.10 Fibre tension	17
3.11 Vibrometer	17
4. DUAL-CORE FIBRE: THEORY	19
4.1 Acousto-optic interaction	20
4.2 Design requirements	22
4.3 Design algorithm	23
4.4 Effect of variations in the <i>a priori</i> parameters	24
5. DUAL-CORE FIBRE: EXPERIMENT	28
5.1 Fibre fabrication	28
5.2 Acoustic transducer	30
5.3 Fibre splicing	32
5.4 Fibre rotational alignment	33
5.5 Fixing of fibre and horn	33
5.6 Acoustic anti-reflection	34
5.7 Conclusion	34

6.	NUL COUPLER FREQUENCY SHIFTER	35
6.1	Description	35
6.2	Advantages over dual-core fibre	36
6.3	Making tapers and couplers	37
6.4	How a taper works	39
6.5	How a coupler works	41
7.	NUL COUPLER: THEORY	43
7.1	The taper transitions	43
7.2	The taper waist	44
7.3	Bandwidth	46
7.4	Polarisation	46
7.5	Longitudinal profile	47
8.	SINGLE TAPER: EXPERIMENT	49
8.1	A single taper	49
8.2	Acoustic wave generation	50
8.3	Acousto-optic resonance	51
8.4	Acoustic reflection	53
8.5	Stability and drift	53
8.6	Parasitic resonances	53
8.7	Another single taper	54
8.8	Acousto-optic resonance	54
8.9	Wavelength dependence	55
8.10	Power requirements	56
8.11	Parasitic resonances	57
9.	NUL COUPLER: EXPERIMENT	58
9.1	Null coupler	58
9.2	Acousto-optic resonance	60
9.3	Wavelength dependence	60
9.4	Drive power	62
9.5	Frequency shift	62
9.6	Repeatability	63
10.	NUL COUPLER USING IDENTICAL STANDARD FIBRES	64
10.1	Null coupler using common fibre	64
10.2	Frequency shifter at 1550 nm using standard single-mode telecom fibre	66
10.3	Frequency shifter at 500 kHz	66
10.3.1	Acousto-optic resonance	66
10.3.2	Wavelength dependence	66
10.3.3	Drive power	66
10.3.4	Frequency shift	67
10.4	Frequency shifter at 12 MHz	70
10.4.1	Acousto-optic resonance	70
10.4.2	Frequency shift	71
10.5	Conclusion	71

By *perform 50*

Distribution	
Availability Codes	
Dist	Avail and/or Special
A-1	

11.	POLARISATION PROPERTIES OF THE SECOND MODES	73
11.1	"Wide" coupler waists	73
11.2	"Narrow" coupler waists	75
11.2.1	Abrupt model	75
11.2.2	Strictly adiabatic model	76
11.3	Experimental	78
11.3.1	Observation of hybrid modes	78
11.3.2	The effect of twist on the null coupler	82
11.4	Conclusion	84
12.	CONCLUSIONS AND PROPOSED FUTURE WORK	85
12.1	Current state of progress	85
12.2	Carrier and image sideband suppression	85
12.3	Robust and stable package	86
12.4	Broadband operation	87
12.5	Polarisation dependence	88
12.6	Low electrical drive power	89
12.7	Higher frequency operation	90
13.	ANCILLARY TOPICS	90
13.1	Publications	90
13.2	Participants	91
13.3	Discoveries, inventions, patents	91
13.4	References	92
14.	APPENDIX: TABLE OF SYMBOLS	93

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By <i>perform 50</i>	
Distribution	
Availability Codes	
DIST	Avail and/or Special
<i>A-1</i>	

1. SUMMARY

The aim of this project is the design and implementation of a practical acousto-optic frequency-shifter compatible with single-mode optical fibre. The frequency shifter was originally to be based on a special fibre with two dissimilar cores. Normally, little light couples between the cores, but in the presence of an appropriate flexural acoustic wave a large transfer of optical power, with a frequency shift, takes place. This interaction has been analysed theoretically, a dual-core fibre frequency shifter designed, and its component parts assembled.

Progress towards this initial experiment was hampered by fibre fabrication problems. In the meantime, an alternative type of frequency shifter was pursued. Here the acousto-optic interaction takes place between two modes in the narrow uniform waist of a special "null" fused taper coupler. This has so many compelling advantages that the earlier dual-core work was put to one side. A prototype null coupler device was built and tested, and immediately satisfied some of the project's goals. The design promises to be able to satisfy the remainder also.

2. INTRODUCTION

This report describes the results of a project to design and make an acousto-optic frequency shifter in single-mode optical fibre. First we list the goals that were identified for the project, and then summarise the two distinct designs that were pursued.

2.1 Project goals

The ideal frequency shifter is an electrically-powered black box with at least an input optical fibre and an output optical fibre. When the device is switched on, all the incoming light in the input fibre is changed in frequency by a given amount (of the order of MHz), and emerges from the output fibre. It is the aim of this project to make such a device. Any real device will deviate from this ideal in a number of ways, and the following specifications have been identified.

Low insertion loss:

Insertion loss is a measure of the amount of input light that fails to emerge as frequency-shifted output. It is defined as

$$\text{insertion loss} = -10 \log_{10} \frac{P_{fs}}{P_{in}} \quad \text{dB ,} \quad (1)$$

where P_{fs} is the frequency-shifted output power, and P_{in} is the input power.

Carrier suppression >50 dB:

In practice, the output frequency-shifted light will be to some extent contaminated by some unshifted light. The unshifted wave is referred to as the *carrier* frequency, using terminology borrowed from the electronic engineer's description of amplitude modulation (the frequency-

shifted wave can be referred to as a *sideband* of the carrier). Carrier suppression is quantified as

$$\text{carrier suppression} = -10 \log_{10} \left[\frac{P_c}{P_{fs}} \right] \quad \text{dB}, \quad (2)$$

where P_c is the power at the carrier frequency.

The output is also usually contaminated by light shifted in the opposite sense to the major frequency-shifted component. For example, if the output is mainly up-shifted, there may also be a small amount of down-shifted light present. This minor component is referred to as the *image sideband*. Its suppression can be quantified in the same way as carrier suppression.

Four- or three-port device:

The ideal frequency shifter described above is a two-port device (one input fibre and one output fibre). The state of the output light when the device is switched off was not specified. It is sometimes desirable to have light routed to a second output fibre when the device is off. This permits a mode of operation in which the device is driven at half-power; half the input light is frequency-shifted and emerges from one output fibre, while the rest is unshifted and emerges from the other output fibre. It is also sometimes an advantage to have a second input fibre, making the device bi-directional. Such a four-port device would be the true fibre analogue of the bulk-optic Bragg cell, and could be incorporated in a number of novel interferometric configurations.

Identical standard single-mode fibre pigtails:

This ensures that the frequency shifter is a convenient component to insert into existing optical fibre networks, which are often constructed out of one standard type of fibre. Compatibility also requires that the device functions at one of the standard optical wavelengths.

Low electrical drive power:

The lower the power required to drive the frequency shifter, the cheaper and more versatile it is. In some remote operating environments, electrical power is at a premium. Low electrical power also permits more compact drive circuitry.

Robust and stable package, no longer than 30 cm:

The device must be capable of withstanding normal use. Again, space is often at a premium in many applications.

Polarisation dependence unspecified:

Many fibre components behave differently depending on the polarisation state of the light passing through. Some designs of frequency shifters may be polarisation insensitive, while others might frequency-shift one polarisation of light while leaving the other polarisation unchanged. Although this clearly affects the way the frequency shifter can be used, no particular polarisation properties have been specified for this project.

High frequency operation not required:

There are frequency shifter applications which require a shift of much greater than a few MHz. However, such large frequency shifts were not specified for this project.

2.2 Realisation of these goals

Two designs of all-fibre frequency shifter were considered, one based on a special dual-core fibre, and the other based on a special type of fused taper coupler. Both designs share the same operating principle as the well-known two-mode fibre frequency shifter. In this device, light in a single-mode input fibre is made to excite just the fundamental mode of a two-mode fibre. A flexural acoustic wave induces coupling to the second mode of the fibre, the coupled light being frequency-shifted. The light in the second mode is then extracted, and emerges from a single-mode output fibre. This design was first publicised some time ago¹. An electrical drive power of 85 mW has been reported².

The two-mode fibre frequency shifter suffers from a number of disadvantages. Firstly, a number of extra optical components must be incorporated, increasing the size and complexity. The performance of the frequency shifter is critically dependent on the performances of these components. A mode filter is necessary at the input, to ensure that no light enters the interaction region in the second mode. Another is necessary to prevent any residual unshifted light in the fundamental mode from contaminating the output. A passive mode convertor is needed to coupler the frequency-shifted wave from the second mode to the fundamental mode, in order to match a single-mode output fibre.

Secondly, being a two-port device, there is no output for unshifted light. In particular, there is no output when the device is in the "off" state.

Thirdly, the design is inflexible, having few degrees of freedom. This makes the design procedure very simple, since the core diameter determines the optical properties and the cladding diameter determines the acoustic properties. However, the resulting fibre is usually highly non-standard, with typically a high core index and a narrow cladding diameter being required. This compromises the fibre's other characteristics, such as ease of fabrication, loss, handability, and compatibility with standard fibre.

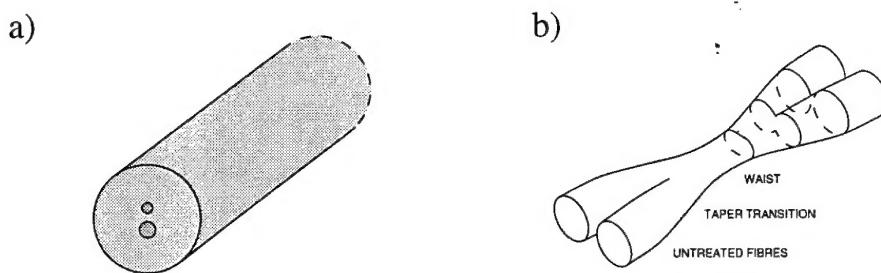


Fig. 1. Schematic dual-core fibre (a) and fused taper coupler (b).

To get around these problems, our original proposal was to study a frequency shifter based on dual-core fibre. This is a special fibre with two parallel single-mode cores at its centre, Fig. 1a. The cores are sufficiently well separated to guide light independently, but are close enough for there to be an optical interaction between them. If the cores are sufficiently dissimilar, the normal modes of the composite structure resemble the modes of the cores in isolation. However, given sufficient modal overlap, resonant coupling between the normal modes can occur in the presence of a flexural acoustic wave. This will result in a transfer of frequency-shifted power from one core to the other.

Use of a dual-core fibre permits the two interacting modes to be spatially separated. The mode filtering and conversion operations of the two-mode device are now achieved by launching and extracting the light at the appropriate core. If launch and extraction can be arranged for both cores at once at each end, a four-port device is possible. The presence of two cores in the fibre gives more freedom in design; the design procedure becomes more complicated, but less compromises are necessary. Dual core fibres with beat periods ranging from a few tens of microns to a few mm can be designed, permitting a wide range of acoustic frequencies of operation.

A dual-core fibre frequency shifter had already been made before the project began³, demonstrating the feasibility of the idea. The first part of the project concentrated on building such a frequency shifter. The fibre and the acousto-optic interactions within it were analysed theoretically, a design procedure worked out, and the component parts of an initial experiment assembled, before work was held up by fibre fabrication problems. There was a delay before new fibre could be supplied.

In the meantime, an alternative and entirely new design of frequency shifter was investigated. This is based on a fused taper coupler, which comprises two input and two output single-mode fibres, linked by a narrow "waist" region formed by fusing and stretching two parallel fibres together, Fig 1b. In a standard coupler, or fibre beam-splitter, light entering one fibre excites two modes in the waist region. The differing phase velocities of these modes, and the distance they travel before the end of the waist region, determines the division of the input light between the two output fibres⁴. An acousto-optic interaction can be made to take place between these modes, but to be usable as a frequency shifter it is necessary to arrange for the input light to excite just one of them. Hence a special type of coupler is required. We have termed this a "null" coupler, because one of its characteristics is that, when passive, it does not actually function as a beam-splitter at all.

It soon became apparent that the new type of device would be superior to the dual-core design, and indeed the two-mode device, for reasons which are detailed in a later section. For example, it would automatically be a four-port device, could be made entirely from standard single-mode fibre, and could have an acoustic power requirement five orders of magnitude lower than the best reported two-mode fibre result. Although no-one had ever made the null fused taper coupler upon which it depended, we felt justified in pursuing this design because of the advantages it offered, and because it conformed to the requirements defined at the start of the project.

The acousto-optic interaction in narrowed fibre waists was demonstrated first of all in a single tapered fibre. Attempts were then made to make a null coupler. This involved the fabrication of couplers from pairs of dissimilar fibres. A suitable pair was found, and one of the first null couplers was made into a prototype frequency shifter. These devices were so successful that work on the dual-core fibre was effectively abandoned.

A large fraction of the project's lifetime was spent on the dual-core fibre work. We have therefore described this work at length, even though it has been superseded. Section 3 describes the theoretical analysis and experimental techniques which are common to both schemes. Sections 4 and 5 describe the dual-core work, and can be omitted if just the "state of the art" is of interest. Sections 6 to 11 describe work leading up to the null coupler device. Remaining sections outline future work in a proposed continuation of the project, and various other topics.

3. ANALYSIS OF THE ACOUSTO-OPTIC INTERACTION IN FIBRES

This section describes theoretical work that is applicable to both designs of frequency shifter. A table of the symbols used in the analysis in this report is given in an appendix (section 14).

3.1 Description

The acousto-optic interaction is fundamentally the same in dual-core fibres and null couplers, and indeed in two-mode fibres. In all cases, an ideal device is constructed so that incoming light in one of two input ports excites one normal mode of an interaction region. In the passive state, light in the excited mode travels along the interaction region to its end unchanged. It then emerges from one of two output ports, which corresponds to the original input port (Fig. 2a). Light entering in the other input fibre excites a second normal mode, travels along the interaction region unchanged, and emerges from the other output port (Fig. 2b). These ports are literally present in the null coupler, correspond to the individual cores in the dual-core fibre, and in the two-mode fibre are provided by whatever machinery is used to spatially separate the two modes. In practice, one or two of the ports may be missing.

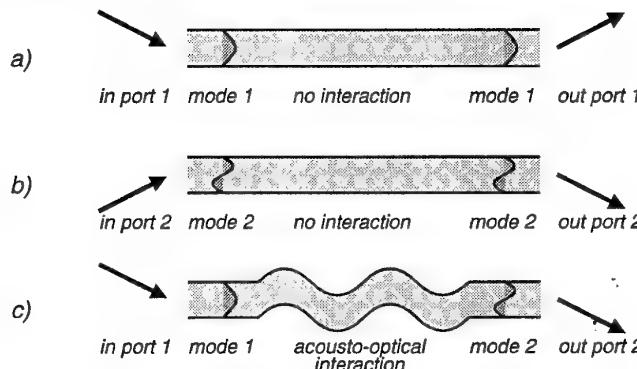


Fig. 2. Behaviour of a fibre acousto-optic device when passive ((a) and (b)) and optimally active (c).

The two normal modes of interest have different phase velocities, and hence different propagation constants. A beatlength L_B can be defined for the two modes. This is the minimum distance two modes, initially in phase, must travel before being once again in phase. If the propagation constants are β_e and β_o , then the beatlength is

$$L_B = \frac{2\pi}{\beta_e - \beta_o}. \quad (3)$$

A flexural acoustic wave imposed on the fibre at the interaction region sinusoidally bends the fibre. This creates an anti-symmetric perturbation of the refractive index distribution across the fibre, which varies periodically along the fibre with a pitch equal to the wavelength Λ of the acoustic wave. If the wave constant $\kappa = 2\pi/\Lambda$ of the acoustic wave matches the difference between the wave constants of the two modes, then the index perturbation causes resonant forward coupling of light between the modes. This acousto-optic resonance condition is the same as saying that the acoustic wavelength matches the optical beatlength:

$$\Lambda = L_B. \quad (4)$$

If the amplitude of the acoustic wave is adjusted so that the light is completely coupled from one mode to the other as the light wave reaches the end of the interaction region, then all of the light will emerge from the other output port, instead of the port from which it emerges when the device is passive (Fig. 2c). Also, because the acoustic wave (and hence the perturbation) travels along the fibre, the coupled light will be frequency-shifted by an amount equal to the acoustic frequency. The cause of the frequency-shift is related to the Doppler shift which occurs when light is reflected by a moving surface.

A plane of symmetry can be defined in each device. This is the plane in which the fibre ports lie, and is also the plane of vibration of the flexural acoustic wave. Cartesian axes are defined so that the z axis points along the fibre, the y axis is transverse to the fibre in the plane of symmetry, and the x axis is transverse to the fibre and normal to the plane of symmetry (Fig. 3).

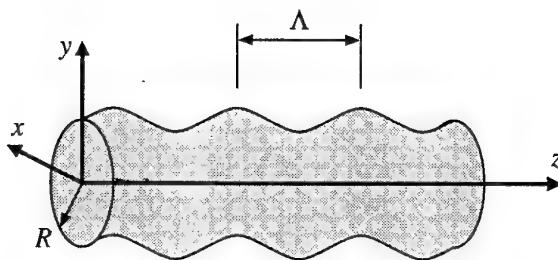


Fig. 3. Cartesian axes for a fibre with a flexural acoustic wave.

The two normal modes of interest are the two lowest-order modes of the interaction region, and will be described as the even (fundamental) and odd (second) modes, although in the asymmetric dual-core fibre the modes have no particular symmetry. In a fibre with near-circular symmetry there are in fact two different second modes distinguished by the spatial orientation of the lobes in their intensity distributions (Fig. 4). In such circumstances, we are interested only in the second mode whose lobes sit on the y axis, so that the field distributions of our even and odd modes are respectively symmetric and anti-symmetric along the y axis.

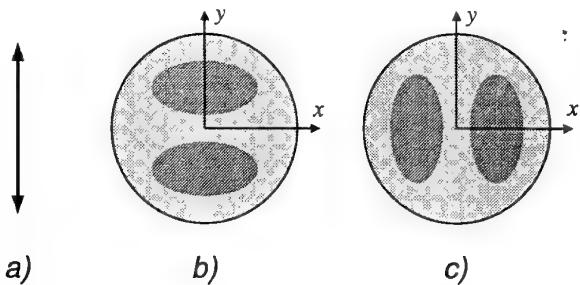


Fig. 4. The cross-section of a near-circularly-symmetric fibre, showing (a) the plane of acoustic vibration, (b) the intensity lobes of the "interesting" second mode, (c) the intensity lobes of the "uninteresting" second mode.

3.2 Passive behaviour

First we consider the output of the device in the absence of an acoustic wave. The acousto-optic interaction as outlined above is idealised, in as much as it is assumed that light in one input fibre excites just one mode of the interaction region. In fact both modes are in general excited, though to approach the ideal one mode should be excited to a much greater extent than the other. Light in the other input fibre excites the orthogonal combination of the same two modes.

The extent to which two modes are excited is made apparent in the passive response of the device.

The field distribution of a light wave entering in ports 1 and 2 is denoted by ψ_1 and ψ_2 respectively, and the field distributions in the even and odd normal modes of the interaction region are denoted by ψ_e and ψ_o respectively. All these are functions of transverse co-ordinates x and y , and are assumed normalised to unity. If a wave entering port 1 excites the normal modes in a power ratio of ξ , the excitation of the normal modes by incoming light can be expressed by

$$\begin{aligned}\psi_1 &= \frac{1}{\sqrt{1 + \xi}} \psi_e - \frac{\xi^{1/2}}{\sqrt{1 + \xi}} \psi_o ; \\ \psi_2 &= \frac{\xi^{1/2}}{\sqrt{1 + \xi}} \psi_e + \frac{1}{\sqrt{1 + \xi}} \psi_o ,\end{aligned}\quad (5)$$

guaranteeing normalisation, orthogonality and power conservation.

The two normal modes excited by an input wave in fibre 1 (say) propagate along the interaction region without coupling, and acquire a phase difference because of their different propagation constants. Hence the field a distance z along the interaction region is

$$\psi(z) = \frac{1}{\sqrt{1 + \xi}} \psi_e e^{-i\beta_e z} - \frac{\xi^{1/2}}{\sqrt{1 + \xi}} \psi_o e^{-i\beta_o z} . \quad (6)$$

The optical powers emerging from the output ports at $z = L$ is found by expressing the final field $\psi(L)$ as a superposition of ψ_1 and ψ_2 , and taking the square modulus of the resulting excitation coefficients. This gives a normalised output power in port 2 (the coupled power) of

$$P_c = M \sin^2 \left(\frac{\pi z}{L_B} \right), \quad (7)$$

with the normalised power in port 1 (the throughput power) being $P_t = 1 - P_c$ as required by power conservation. The maximum intrinsic coupling M is given by

$$M = \frac{4\xi}{(1 + \xi)^2} . \quad (8)$$

As L or L_B are varied (L_B can be varied by changing the wavelength of the input light), the passively coupled power P_c can vary between zero and M , as a fraction of the input power. Thus M can be measured experimentally. The excitation ratio ξ can then be calculated using (8).

Performing the same analysis for incoming light in input port 2 gives identical expressions, but with the coupled light now emerging from output port 1.

3.3 Acousto-optic interaction

The optical consequences of the presence of a flexural acoustic wave are now considered. The frequency f of the acoustic wave is known, being given by the frequency at which the acoustic generator is driven. The acoustic wavelength is given in terms of f by the dispersion relation for flexural acoustic waves on a cylinder. This relation is generally complicated, but in the low frequency limit can be approximated to⁵

$$\Lambda \equiv \frac{2\pi}{\kappa} = \left[\frac{\pi R c_{ext}}{f} \right]^{\frac{1}{2}}, \quad (9)$$

where R is the fibre radius and c_{ext} is the speed of extensional waves in silica and has the value 5760 m s^{-1} .

The low frequency limit holds if $\Lambda \gg 2R$. As well as complicating the dispersion relation, higher frequencies cause practical problems as well. The strain distribution across the fibre ceases to be simple, and the wave begins to take on the character of a surface wave. This weakens the acousto-optic interaction because the light is usually guided at the centre of the fibre. Hence only acoustic waves with wavelengths large compared to the fibre diameter are considered. This restricts the analysis to relatively low frequencies of a few MHz or less for fibres with conventional diameters of the order of $100 \mu\text{m}$, though for the narrow waists in a null coupler the upper limit is several hundred MHz.

The low frequency acoustic wave periodically bends the fibre. This affects the refractive index distribution across the fibre in two ways⁶. There is a true change in the index of the material due to the elasto-optic effect, which reduces the index locally in proportion to the strain. There is also an effective change in the index due to the change in optical path caused by the bending, which increases the index locally in proportion to the strain. In fact the latter geometrical effect dominates the material effect. The net change Δn in the index n_o is given by

$$\Delta n = n_o(1 + \chi)S_z, \quad (10)$$

where S_z is the longitudinal strain and χ is related to the elasto-optic constants. For silica, $\chi = -0.22$. The strain distribution within the fibre is given by

$$S_z(x, y, z) = -y \frac{d^2 u(z)}{dz^2}, \quad (11)$$

where $u(z)$ is the fibre displacement caused by the acoustic wave. Hence for a sinusoidal acoustic wave of amplitude u_o and angular frequency $\Omega = 2\pi f$,

$$u(z) = u_o \cos(\Omega t - \kappa z), \quad (12)$$

there is a travelling variation of the effective refractive index along the fibre:

$$\begin{aligned} \Delta n(x, y, z, t) &= n_o(1 + \chi)\kappa^2 u_o y \cos(\Omega t - \kappa z) \\ &\equiv \Delta n(x, y) \cos(\Omega t - \kappa z). \end{aligned} \quad (13)$$

This change in the index distribution perturbs the optical modes travelling along the fibre and causes coupling between them.

In passing, we note the expression for the power carried by a flexural wave of amplitude u_o :

$$P = 4\rho(\pi^7 c_{ext} R^5 f^5)^{\frac{1}{2}} u_o^2, \quad (14)$$

where the density ρ of silica is 2200 kg m^{-3} . This expression makes use of the low frequency dispersion relation (9).

3.4 Frequency shifter output powers

Monochromatic light is launched into one input port only, and a travelling flexural acoustic wave of given amplitude and frequency is excited on the fibre. The object of the analysis is to determine the powers and frequency spectra of the light in either output port.

The acoustic wave gives rise to a travelling periodic variation of the effective refractive index along the fibre. This couples together the two normal modes, which therefore cease to be invariant waves. The invariant waves are now Bloch modes of the structure, which are solutions of the wave equation with the same periodicity as the index variation. An "average propagation constant" $\bar{\beta}$ and an acousto-optic mismatch δ can be defined:

$$\bar{\beta} = \frac{1}{2}(\beta_e + \beta_o + \kappa) ; \quad (15a)$$

$$\begin{aligned} \delta &= \frac{1}{2}(\beta_e - \beta_o - \kappa) \\ &= \pi \left(\frac{1}{L_B} - \frac{1}{\Lambda} \right), \end{aligned} \quad (15b)$$

so that δ is zero at acousto-optic resonance (4). Introducing the angular frequency ω of the light, the propagation constants β_{\pm} and the field distribution functions $\Psi_{\pm}(x, y, z, t)$ of the two Bloch modes can be written as

$$\begin{aligned} \beta_{\pm} &= \bar{\beta} \pm \sqrt{\delta^2 + C^2} ; \\ \Psi_+ &= e^{i(\omega t - \beta_{+}z)} [\psi_e \cos(\phi) + \psi_o \sin(\phi) e^{-i(\Omega t - \kappa z)}] ; \\ \Psi_- &= e^{i(\omega t - \beta_{-}z)} [-\psi_e \sin(\phi) + \psi_o \cos(\phi) e^{-i(\Omega t - \kappa z)}], \end{aligned} \quad (16)$$

where $\tan(2\phi) = C/\delta$ so that ϕ is in the first quadrant ($= \pi/4$ on resonance), and C is the acousto-optic coupling coefficient defined by

$$C = \frac{k}{2} \int \int_A \psi_e(x, y) \psi_o(x, y) \Delta n(x, y) dx dy, \quad (17)$$

with $\Delta n(x, y)$ from (13). k is $2\pi/\lambda$, where λ is the free space optical wavelength. The above expressions for the Bloch modes are valid provided that both δ and C are small compared to κ . The integral for C cannot be evaluated until the particular design of frequency shifter is decided upon. However, we can immediately see that it is proportional to the acoustic amplitude u_o . It is also clear that there can be no coupling between two modes whose y dependences are either both even or both odd, because $\Delta n(x, y)$ is itself an odd function of y . This is why in a fibre with near-circular symmetry we only consider that second mode whose lobes lie on the y axis, since the other has an even y dependence.

The output light waves from the device are determined by expressing the input wave as a superposition of the Bloch modes, and examining the form of the superposition at the output end. As expected, the light emerging from the far end of the fibre is a mixture of three optical frequencies - the carrier (i.e. the input frequency), and two sidebands displaced up and down from the carrier by the acoustic frequency.

For a launched light wave in input port 1, the normalised output powers at $z = L$ are

$$P_1(0) = \left[1 - \frac{C^2}{\delta^2 + C^2} \sin^2(L\sqrt{\delta^2 + C^2}) \right] \left[1 - M \sin^2\left(\frac{\kappa L}{2} + \varepsilon(L)\right) \right]; \quad (18a)$$

$$P_1(+f) = \frac{C^2}{\delta^2 + C^2} \frac{M}{4} \sin^2(L\sqrt{\delta^2 + C^2}); \quad (18b)$$

$$P_1(-f) = P_1(+f); \quad (18c)$$

$$P_2(0) = \left[1 - \frac{C^2}{\delta^2 + C^2} \sin^2(L\sqrt{\delta^2 + C^2}) \right] M \sin^2\left(\frac{\kappa L}{2} + \varepsilon(L)\right); \quad (18d)$$

$$P_2(+f) = \frac{C^2}{\delta^2 + C^2} \frac{[1 \mp \sqrt{1 - M}]^2}{4} \sin^2(L\sqrt{\delta^2 + C^2}); \quad (18e)$$

$$P_2(-f) = \frac{C^2}{\delta^2 + C^2} \frac{[1 \pm \sqrt{1 - M}]^2}{4} \sin^2(L\sqrt{\delta^2 + C^2}), \quad (18f)$$

where $P_i(\Delta\nu)$ is the optical power leaving output port i with a frequency shift of $\Delta\nu$. In the expressions for P_2 , the upper sign holds if $\xi < 1$ (that is, the input light excites the even mode preferentially), whereas the lower sign holds if $\xi > 1$. $\varepsilon(L)$ is a slowly-varying and somewhat uninteresting function, given for completeness by

$$\tan^2 \varepsilon = \frac{\delta^2}{\delta^2 + C^2} \tan^2(L\sqrt{\delta^2 + C^2}). \quad (19)$$

M is the maximum intrinsic coupling of the passive device, given by (8).

The set of equations (18) give the output powers in both ports at the carrier frequency and at both sidebands. These powers depend on the acoustic wave amplitude through C in (17), and on the acoustic frequency through δ in (15b).

The parameters C , L_B and M remain to be determined. To calculate them, the particular form of the device must first be specified. However, we can now express various useful quantities such as the acoustic amplitude and power for maximum conversion, carrier and image sideband suppression, optical and acoustic bandwidth, and polarisation dependence, in terms of the three unknowns.

3.5 Maximum conversion

The theoretical conversion efficiency from light in input port 1 to frequency-shifted light in output port 2 is given by either $P_2(+f)$ or $P_2(-f)$ in (18e) and (18f), depending on whether the device is configured as an up-shifter or a down-shifter. The maximum conversion takes place when $M = 0$, $\delta = 0$, and

$$CL = \frac{\pi}{2}. \quad (20)$$

M is intrinsic to the particular device and is determined when it is constructed. For small M , the maximum acousto-optic conversion efficiency is $1 - M/2$.

The condition on δ is simply the condition for acousto-optic resonance (4), and is satisfied for particular values of f and λ in a given device:

$$L_B(\lambda) = \Lambda(f). \quad (21)$$

The function $\Lambda(f)$ is known immediately through (9), but $L_B(\lambda)$ depends on the detailed optical properties of the particular fibre on which the device is based.

The condition (20) determines an optimum value for the amplitude of the acoustic wave, through (17) with $\Delta n(x, y)$ from (13) and Λ given by (9),

$$u_o = \frac{1}{8\pi} \frac{c_{ext}}{(1 + \chi)n_o} \frac{R\lambda}{Lf} \frac{1}{I_o}, \quad (22)$$

and hence an optimum value for the acoustic power through (14). The overlap integral I_o ,

$$I_O = \iint_{A_\infty} \psi_e \psi_o y dx dy, \quad (23)$$

contains information about the fibre's structure and optical properties and, via the constraint provided by the resonance condition (21), may also provide some further f dependence in (22).

We can immediately see that the required acoustic amplitude is inversely proportional to the interaction length L , and so the acoustic power is inversely proportional to L^2 . Shorter devices require a greater acoustic power.

3.6 Frequency shift

The magnitude of the optical frequency shift $\Delta\nu$ is clearly equal to the acoustic frequency,

$$\Delta\nu = \pm f, \quad (24)$$

but the sign of the frequency shift depends on the directions of the acoustic and optical waves along the coupler, and also on which of the even mode and the odd mode is excited by the input light. In fact, since for non-zero M both modes are excited to some extent, there will be output light with both signs of frequency shift; in (18), both $P_2(+f)$ and $P_2(-f)$ are non-zero. However, when M is small, one component will be very much smaller than the other, so we can treat the smaller component as a contaminant in the main frequency-shifted wave.

From (18e) and (18f) with a forward-travelling acoustic wave given by (12), we find that the frequency shift is positive (up-shift) if the input light excites mainly the odd mode, and negative (down-shift) if the input light excites mainly the even mode. If the acoustic wave is backward-travelling relative to the light wave, we can replace Ω with $-\Omega$ in (12), and hence f with $-f$ in (18), giving the opposite frequency shift. Hence in a given device, interchanging the input ports, or swapping the input ports with the corresponding output ports (equivalent to reversing the acoustic wave), reverses the sign of the frequency shift.

	launch in even mode	: launch in odd mode
co-propagating acoustic wave	down-shift	up-shift
counter-propagating acoustic wave	up-shift	down-shift

Table I. The acousto-optic frequency shift in various circumstances.

Clearly, if one case can be remembered, then the others follow by changing sign the appropriate number of times. There is a simple mnemonic for the case where the even mode is launched and the acoustic wave is co-propagating, illustrated in Fig. 5. Consider a plane wave propagating forwards in free space; it has a positive propagation constant. The wave encounters a plane mirror at normal incidence and is reflected. Now it has a negative propagation constant, so the propagation constant is reduced. If the mirror is moving in the same direction as the input wave, the reflected wave is reduced in frequency by the Doppler effect. The situation is analogous to the acousto-optic device: the scattering surface is the forward-moving index modulation associated with the acoustic wave; and the propagation constant is reduced as the light is coupled from the even mode (with the greater propagation constant) to the odd mode (with the lesser propagation constant).

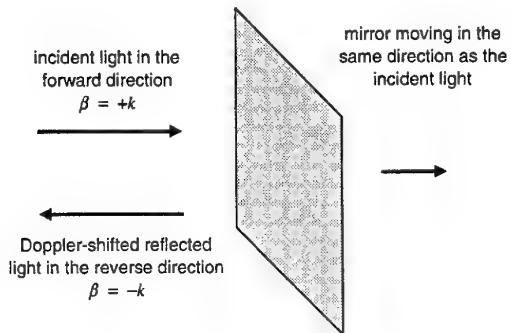


Fig. 5. A mnemonic for determining the sign of a frequency shift. The frequency of a light wave is down-shifted, by the Doppler effect, upon reflection at a forward-moving mirror.

3.7 Bandwidth

The various designs of frequency shifter can in general be operated over wide ranges of acoustic frequency and optical wavelength. However, if (say) the frequency is changed, the wavelength of operation also changes to compensate, in accordance with the resonance condition (21). Here we consider "bandwidth" to mean the range of acoustic frequencies the device can operate over once a given optical wavelength has been chosen, and *vice versa*. In fact we can generalise the concept of bandwidth and consider the effects of a variation of any parameter q on conversion efficiency. The range of q for which the conversion efficiency is over half the maximum value is termed the FWHM bandwidth Δq , Fig. 6.

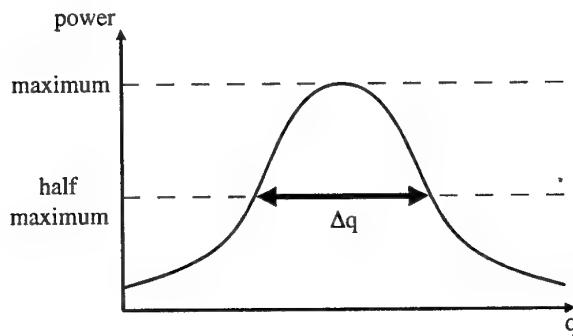


Fig. 6. The full width Δq at half maximum of a resonance as q varies.

The frequency-shifted power is given by $P_2(+f)$ or $P_2(-f)$ from (18). The deviation from resonance as q is varied is given by $\delta(q)$ from (15b). The δ dependence of the sinusoidal parts of P_2 complicates the bandwidth calculation, but a good approximation is obtained by ignoring this and considering just the rational factor. This reduces to half its maximum value when $\delta = C$, so the FWHM bandwidth in terms of δ is $\Delta\delta = 2C$. (A more thorough calculation, incorporating the sinusoidal dependence on δ in (18), results in the replacement of the factor 2 with 1.6.)

The value of C is given simply from the result that, at the optimum operating condition, (20) holds. The translation from the somewhat abstract quantity $\Delta\delta$ to the interesting bandwidth Δq results from considering the differential dependence of δ on q ; (15b) and the chain rule for partial derivatives gives the general result for small Δq

$$|m_q| \frac{\Delta q}{q} = \frac{2\Lambda}{L}, \quad (25)$$

where

$$m_q = -2q \frac{\partial}{\partial q} \left(\frac{\Lambda}{L_B} \right) \quad (26)$$

evaluated at $\Lambda/L_B = 1$. This gives a bandwidth which is simply related to the number of acoustic wavelengths (or optical beatlengths) which fit in the interaction length.

The bandwidths of specific quantities such as λ depend on the particular form of $L_B(q)$. However, Δf is a special case, since L_B is never a function of f and $\Lambda(f)$ is known from (9). This gives constant $m_f = 1$. Indeed, f illustrates a general rule that if Λ/L_B always varies as a given power of q about the resonance condition, m_q can be found by writing this variation as $\sim q^{-m_q/2}$.

3.8 Compensation for one parameter by another

Similar considerations give the shift Δq_1 in the resonance value of the parameter q_1 , caused by a variation Δq_2 in another parameter q_2 :

$$m_1 \frac{\Delta q_1}{q_1} = m_2 \frac{\Delta q_2}{q_2}. \quad (27)$$

3.9 Carrier and image sideband suppression

Carrier suppression refers to the extent to which the frequency-shifted output is contaminated by unshifted light. The amount of unshifted light is given by $P_2(0)$ in (18d). On resonance at maximum conversion, this is identically zero. At the edge of the resonance band, where $\delta = \pm C$, the unshifted power is approximately given by

$$P_2(0) \sim \frac{M}{4}. \quad (28)$$

Carrier suppression is usually expressed in decibels relative to the amount of frequency shifted light. Since we can assume that the latter is near unity:

$$\text{carrier suppression} = -10 \log_{10} (P_2(0)). \quad (29)$$

Image sideband suppression also refers to the contamination of the frequency-shifted output light, this time by a component with the opposite shift. As indicated above, this is due to the excitation of both modes at the input, to some extent. The amount of power in the main frequency shifted wave at the output is given by the greater of $P_2(+f)$ and $P_2(-f)$ in (18); that is, whichever contains the positive root. The lesser component, with the negative root, gives the power in the image sideband. For small M , this is

$$P_2(\text{image sideband}) = \frac{M^2}{16}. \quad (30)$$

A second source of image sideband contamination is the presence of a backward-travelling acoustic wave in the interaction region (in this context, "backward" is relative to the direction of propagation of the major part of the acoustic wave). In some schemes, the acoustic wave is imposed upon the fibre in the interaction region, so some directly-generated backward component is inevitable. In both the dual-core and null coupler designs considered in this

project, however, we avoid any directly-generated backward wave. The remaining source of backward acoustic waves is acoustic reflection beyond the interaction region. This directs a proportion of the incident acoustic wave back through the interaction region. Since a reversed acoustic wave generates a reverse frequency-shift, the resulting image sideband power is given by the same expression for $P_2(\pm f)$, but with a reverse coupling coefficient for C . Since C is proportional to the square root of the acoustic power through (17) and (14), and (20) holds for the main acoustic wave at maximum conversion, the image sideband power for small values of acoustic power reflection coefficient R_{ac} is

$$\begin{aligned} P_2(\text{image sideband}) &= \sin^2\left(\frac{\pi}{2}\sqrt{R_{ac}}\right) \\ &\approx \frac{\pi^2 R_{ac}}{4}. \end{aligned} \quad (31)$$

Again, the image sideband suppression is usually expressed in decibels:

$$\text{image sideband suppression} = -10 \log_{10} (P_2(\text{image sideband})). \quad (32)$$

3.10 Fibre tension

The acoustic velocity is to some extent dependent on the tension in the fibre. Hence the acoustic dispersion relation (9) should include corrections for tension. It can be shown that the new dispersion relation is

$$\frac{\Lambda^2}{\left(1 + \frac{s\Lambda^2}{\pi^2 R^2}\right)^{1/2}} = \frac{\pi R c_{ext}}{f}, \quad (33)$$

the denominator on the LHS containing the correction for the fibre strain s . From this expression, the change $\Delta\Lambda$ in the acoustic wavelength brought about by a small strain can be written as

$$\frac{\Delta\Lambda}{\Lambda} = \frac{c_{ext}}{4\pi R f} s. \quad (34)$$

The resulting shift in the value of parameters such as f and λ can be calculated from (27), given that $m_\Lambda \equiv -2$.

Although we show experimentally that fibre tension can make a large difference to the resonance wavelength, we cannot yet measure the strain imposed upon the fibre. For now, we simply note that the effect of tension can be analysed, and otherwise ignore it.

3.11 Vibrometer

In a full characterisation of the action of a frequency shifter, and indeed of the transducer assembly which generates the acoustic wave, it is necessary to determine the amplitude of acoustic vibration at a surface, be it a fibre, a horn tip or the face of a PZT disc. Given the existing expertise of members of the ORC working in the field of fibre sensors, a non-contact vibrometer has been constructed in situ using a single-mode optical fibre technique.

The optical side of the vibrometer is based on a directional fibre coupler, Fig. 7. Light is launched into one of the input ports of the coupler from a helium-neon laser, and the light

emerging from the second "input" port is detected by a fast detector circuit constructed for the vibrometer. The coupler has a high directivity, so the light detected is almost entirely that reflected back into the coupler from the "output" fibres. Reflection from one of these output fibres is eliminated by immersing the end in a capsule of index-matching fluid. Reflection from the other ordinarily comes from Fresnel reflection at its end, with about 4 % of incoming light being returned. However, if a second reflecting surface is positioned close to the end of the fibre, such that a substantial proportion of the reflected light is recaptured by the fibre, the light returned in the fibre will be the product of interference between the external and Fresnel reflected waves. In other words, the external surface and the fibre end, separated by air, form a very low finesse Fabry-Perot interferometer, in which only two orders of reflection are important.

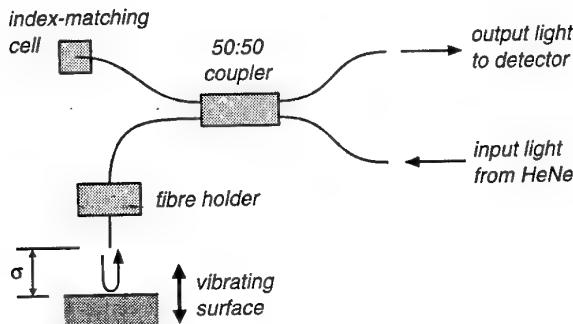


Fig. 7. The optical fibre vibrometer.

If the external surface is vibrating, then the returned light will be modulated at the frequency of vibration, and with a modulation spectrum given by the amplitude of the vibration. It is this spectrum which is measured, and from which the character of the vibration is found. In fact all the fibre coupler usefully does is extract this reflected light wave from the fibre.

If the separation of the reflecting surfaces is σ , the amplitude reflection coefficient of the fibre endface is r_1 , and the normalised amplitude of the wave externally reflected and recaptured by the fibre is r_2 , then the normalised amplitude E which returns to the coupler is

$$E = r_1 + r_2 e^{2ik\sigma}, \quad (35)$$

and the fraction of injected light power which is detected, after two passes through the 50:50 coupler, is

$$P_o = \frac{1}{4} [|r_1|^2 + |r_2|^2 + 2|r_1||r_2| \cos(2k\sigma + \eta)], \quad (36)$$

where η is a constant given by the phase changes on reflection at each surface. If the separation σ is time-dependent because the external surface is vibrating with frequency f and amplitude u_o , so that

$$\sigma = \sigma_o + u_o \sin(2\pi ft), \quad (37)$$

then the detected signal will have a spectrum of frequencies given by

$$V_{detected} \propto \cos(2k\sigma_o + \eta) [J_0(ku_o) + 2J_2(ku_o) \cos(2 \times 2\pi ft) + \dots] - \sin(2k\sigma_o + \eta) [2J_1(ku_o) \sin(2\pi ft) + 2J_3(ku_o) \sin(3 \times 2\pi ft) + \dots]. \quad (38)$$

Hence from a knowledge of the electrical powers in the first and third harmonics of the detected signal, or the peak powers in the first and second harmonics if the mean separation σ_o is itself slowly scanned across an interferometer fringe, the amplitude of the vibration can be calculated.

The sensing fibre tip is carefully cleaved and mounted on a piezoelectric rod. By supplying a voltage to the rod, the position of the sensing tip can be finely altered, or varied sinusoidally at a low frequency if that should prove useful. The rod is mounted on an XYZ positioner to locate the tip over the appropriate surface to be measured. It is possible to position the fibre end close enough to the moving surface to achieve almost complete visibility of the interferometer fringes. The frequency spectrum of the vibrometer's optical output is measured with an RF spectrum analyser. An example of a typical spectrum is given in Fig. 8; the moving surface here was a large piezoelectric cylinder driven at about 19 kHz. The graph displays the peak-holding spectrum as σ_o was tracked across a complete interferometer fringe, so the ratio of first to second harmonics, which was 3.75, is sufficient to give the amplitude of the vibration, in this case $u_o \approx 51$ nm.

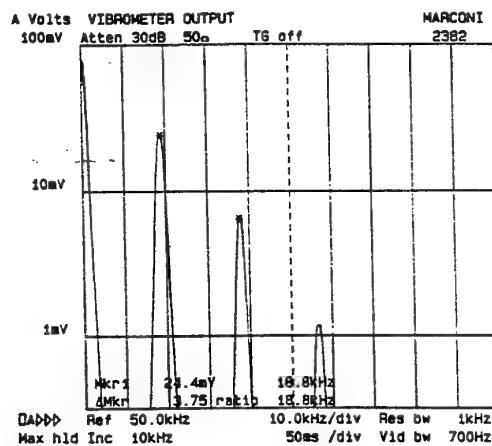


Fig. 8. The output of an RF spectrum analyser at the vibrometer output.

4. DUAL-CORE FIBRE: THEORY

As a first step in the development of a dual-core fibre frequency shifter, the construction of the initial experimental device was to be based on well-known technology, with an emphasis on simplicity rather than performance. This would allow us to gain experience with the problem, and still permit the incorporation of novel features. Once a frequency shifter is demonstrated, it would be possible to improve the performance of each aspect of the device incrementally. Accordingly, the familiar horn transducer configuration was chosen as the starting point. Construction and operation was to be as follows:

A length of dual-core fibre, about 30 cm of which has been stripped of its coating, is fusion-spliced to a single-core fibre, so that the single-core fibre delivers light into only one of the cores of the dual-core fibre. The fibres are held so that one core in the dual-core fibre lies vertically above the other. The transducer/horn assembly is positioned under the single-core fibre near the splice, and the tip of the horn is fused to the fibre. The PZT disc at the base of the horn is driven by an RF signal with the desired frequency. The longitudinal acoustic wave

produced is concentrated by the tapering horn, exciting a flexural wave on the fibre which travels away from the horn in both directions. Mechanical damping where the fibres are held minimises acoustic reflection; this, together with the positioning of the horn at the single-core fibre, ensures that the acoustic wave is unidirectional along the dual-core fibre. The amplitude of the acoustic wave on the fibre is measured using a non-contact vibrometer. The vibrometer is constructed *in situ* and is described in section 3.11.

In the absence of acoustic excitation, light launched into one core of the dual-core fibre remains substantially in that core along the length of the fibre because the intrinsic coupling is small. If an acoustic wave with a wavelength matching the beatlength of the fibre is present, resonant coupling between the two normal modes takes place. Given an appropriate acoustic amplitude and device length, a frequency-shifted light wave is efficiently coupled into the second core, from which it can be extracted by an imaging system or a second splice to a single-core fibre.

As part of the drive towards simplicity in this experiment, the design was based on a fixed optical wavelength of 633 nm. This permits the use of a helium-neon laser as the light source, which has many advantages over the alternatives. The visibility of the light greatly simplifies setting up, the light is narrow-band, and optical components are readily available for this wavelength. This fixing a priori of the optical wavelength contrasts with the preliminary experiment conducted before the start of the project³, where a tunable light source (white light and monochromator) was used. There, the acoustic transducer could be driven at its natural frequency to maximise amplitude, and the optical wavelength subsequently adjusted to achieve acousto-optic resonance. This degree of freedom is not available in the new experiment, and instead the acoustic frequency would be adjusted to achieve acousto-optic resonance. Thus the transducer would be driven to some extent away from its natural frequency, reducing the attainable acoustic amplitude. However, this is closer to the context in which a frequency shifter would be designed in practice, and so is a problem which would have to be overcome at some stage anyway.

A range of relatively low acoustic frequencies, around 1.5 MHz, were chosen for the experiment. This simplifies the analysis (and hence the design) of the acoustic properties of the fibre, because the acoustic displacement can be assumed to be simply distributed across the fibre cross-section (this is the low-frequency limit for which (9) is valid). If higher frequencies were chosen, such that the acoustic wavelength on the fibre is not large compared to the fibre diameter, then acoustic power would concentrate on the fibre surface away from the cores, not only complicating the analysis but also reducing the efficiency of the device.

The design of the elements of this experiment is now described in more detail, as are the techniques for their assembly. The design and construction of each part of the experiment had been mastered before work stopped, due to difficulties experienced with the fabrication of the dual-core fibre. It was therefore not possible to assemble and test a complete device.

4.1 Acousto-optic interaction

We consider a weakly-guiding two-core fibre with the cores sufficiently mismatched in size for the intrinsic coupling M to be small, and sufficiently well-separated (separation d) for coupling to be weak. The structure is depicted in Fig. 9. The normal modes which acousto-optically interact are the so-called supermodes of the structure. Because coupling between the cores is weak, each supermode is a superposition of the fundamental modes of the two cores in

isolation. Because the fibre is not symmetric along the y axis, the normal modes of the fibre are in fact neither even nor odd. However, one is quasi-even in the sense that the field amplitude has the same sign in each core, while the other is quasi-odd with the field amplitude having opposite signs in the two cores.

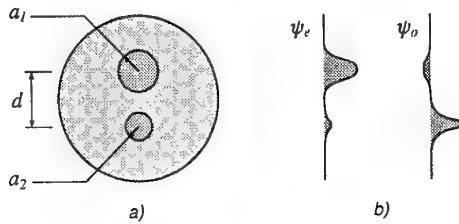


Fig. 9. (a) The cross-section of a dual-core fibre. (b) The quasi-even ψ_e and quasi-odd ψ_o normal modes of a dual-core fibre.

To apply the general analysis above to this specific case, it is necessary to find expressions for the three quantities L_B (giving the resonance condition), I_O (giving the acoustic amplitude and power for maximum conversion) and M (giving the carrier and image sideband suppression, and also contributing to I_O).

L_B is given by the propagation constants of the quasi-even and -odd modes. Since M is small, then these are well approximated by the propagation constants β_1 and β_2 of the two cores in isolation. The beatlength can then be written as

$$\begin{aligned} L_B &= \frac{2\pi}{\beta_1 - \beta_2} \\ &= \frac{2 n_o \lambda}{(NA)^2} \left[\frac{W_1^2}{V_1^2} - \frac{W_2^2}{V_2^2} \right]^{-1}, \end{aligned} \quad (39)$$

where

$$V = ka(NA) \quad (40)$$

is the normalised frequency of a core, the universal function

$$W(V) = a \sqrt{\beta^2 - k^2 n_{cladding}^2} \quad (41)$$

is the well-known eigenvalue for the fundamental mode of a circular fibre, the numerical aperture

$$NA = \sqrt{n_{core}^2 - n_{cladding}^2}, \quad (42)$$

and $n_o = 1.458$ is the refractive index of silica.

The overlap integral I_O is given by (23). Since the two cores provide the input and output ports of the device, the normal mode fields ψ_e and ψ_o are given in terms of the fields ψ_1 and ψ_2 of the cores in isolation by inverting (5). Assuming that the core fields are well-localised at each core, so that the fields are only significant when $y = \pm d/2$, it can be shown that

$$I_O = M^{1/2} \frac{d}{2} \quad (43)$$

with the aid of (8). Hence the acoustic amplitude for maximum conversion is given from (22) by

$$u_o = \frac{1}{4\pi} \frac{c_{ext}}{(1+\chi) n_o} \frac{R\lambda}{Lf d} \frac{1}{M^{1/2}}, \quad (44)$$

with the corresponding acoustic power from (14).

The maximum intrinsic coupling M of the dual-core fibre is determined by the fibre's core geometry and index, and is given by

$$\frac{1}{M} = 1 + \frac{X_d^2}{4} \left[\frac{W_1^2}{V_f^2} - \frac{W_2^2}{W_2^2} \right]^2, \quad (45)$$

where X_d is a complicated but analytic function, which contains the dependence on the core separation d in the first factor displayed:

$$X_d^2 = \frac{V_1^2 V_2^2}{4 K_0 \left(W_1 \frac{d}{a_1} \right) K_0 \left(W_2 \frac{d}{a_2} \right)} \frac{K_1^2(W_1) K_1^2(W_2)}{K_0(W_1) K_0(W_2)} \left[1 + \frac{V_1^2 W_2^2}{U_1^2 V_2^2} \right] \left[1 + \frac{V_2^2 W_1^2}{U_2^2 V_1^2} \right]^{-1} \\ \times \left\{ \frac{1}{K_0 \left(\frac{V_1}{V_2} W_2 \right)} + I_0 \left(\frac{V_1}{V_2} W_2 \right) \left[W_2 \frac{K_1(W_2)}{K_0(W_2)} - \frac{V_1}{V_2} W_2 \frac{K_1 \left(\frac{V_1}{V_2} W_2 \right)}{K_0 \left(\frac{V_1}{V_2} W_2 \right)} \right] \right\}^{-1} \\ \times \left\{ \frac{1}{K_0 \left(\frac{V_2}{V_1} W_1 \right)} + I_0 \left(\frac{V_2}{V_1} W_1 \right) \left[W_1 \frac{K_1(W_1)}{K_0(W_1)} - \frac{V_2}{V_1} W_1 \frac{K_1 \left(\frac{V_2}{V_1} W_1 \right)}{K_0 \left(\frac{V_2}{V_1} W_1 \right)} \right] \right\}^{-1}. \quad (46)$$

U is given by $U^2 = V^2 - W^2$, I and K are modified Bessel functions, and a_1 and a_2 are the radii of cores 1 and 2. The expression for X_d is cumbersome, but its simpler form for identical cores ($V_1 = V_2$) deviates significantly and would make a bad approximation.

4.2 Design requirements

It is now possible to predict the performance of a dual-core fibre frequency shifter, given its structural and operational parameters and the assumptions made above. The relevant requirements for the design of the fibre are as follows:

- a) The fibre has sensible and realistic properties as a fibre. That is, the fibre should be handleable and strong, its dimensions and refractive indices should be within the usual ranges for fibre fabrication, both cores should be single-moded but resistant to bend loss at the wavelength of interest, and sufficiently well separated to be individually selectable for light input and output. A typical fibre has a diameter of the order of 100 μm , a pure silica cladding, a germania-doped core giving an NA between 0.10 and 0.20, and a core V value of around 2.0. The choice of optical wavelength and V fixes the LP_{11} cut-off wavelength of the core to be

$$\lambda_{cutoff} = \frac{V\lambda}{2.405}. \quad (47)$$

- b) The overall interaction length L of the device does not exceed about 30 cm. The acoustic amplitude necessary to drive the device should be minimised, given this length constraint.

Provided the maximum available acoustic amplitude exceeds that required in (44) and overcouples the device, the condition can always be satisfied by reducing the input to the transducer. This equation indicates that a large maximum intrinsic coupling M is desirable in this respect (although M must be small if the quality of the frequency-shifted output is not to be compromised).

- c) The acoustic frequency chosen satisfies the low-frequency condition $\Lambda \gg 2R$, which amounts to

$$f \ll \frac{c_{ext}}{4\pi R} \quad (48)$$

given the dispersion relation of (9). The right-hand side is about 9 MHz for a 100 μm diameter fibre.

- d) The device is operated at acousto-optic resonance, so that (21) holds for the optical wavelength and acoustic frequency of interest. However, unavoidable variations from the intended design (inevitably introduced during fibre fabrication) mean that some tuning of the fibre properties, optical wavelength or acoustic frequency is necessary in practice to obtain resonance. A limited range of fibre properties is possible by making a family of slightly different fibres, but the final process of tuning to resonance is through the acoustic frequency, the optical wavelength being fixed at 633 nm. The fibre tension could also be used for fine adjustment.
- e) In an ideal frequency shifter, all of the input optical power is converted from the initial frequency to only one of the two shifted frequencies. Hence both the unshifted (carrier) and the other shifted (image sideband) waves should be suppressed. Equations (28) and (30) both indicate that we require M to be as small as possible. This must be compromised with the requirement that the device has a low drive power, with (44) indicating a large M .
- f) A further requirement for efficient suppression of unwanted frequencies is that the acoustic wave on the fibre be as unidirectional as possible. Any acoustic power propagating in the other direction will strive to frequency-shift the light in the opposite sense to that desired, increasing the image sideband signal. Hence the primary acoustic wave is excited in the single-core fibre (so that it travels in one direction only along the dual-core fibre), and the fibre holding arrangement absorbs as much acoustic power as possible to minimise secondary reflected waves.

4.3 Design algorithm

A procedure is identified which, once values for the optical wavelength λ , acoustic frequency f , maximum intrinsic coupling M , fibre cladding diameter $2R$, core V values V_1 and V_2 and device length L have been chosen, determines the values of all other parameters necessary to ensure efficient frequency shifter operation. The first three of these input parameters are fixed by the required specification of the eventual device (the target M is determined via (30) by the required image sideband suppression ratio of the shifter), while the rest are chosen for convenience.

The fibre radius and acoustic frequency give Λ via (9), and L_B is given by (39). For efficient frequency shifting, the acoustic and optical beatlengths must match. The wavelength and convenient V values are chosen beforehand, so beatlength matching fixes NA . Each core radius

a_1 and a_2 is then determined from the definition of V .

M is determined by the above parameters and by the centre-centre separation M of the fibre cores, using (45), with d chosen so that the desired value of M results. The acoustic amplitude u_0 required for maximum conversion is then given by (44).

For the particular fibre which was to be used in the initial experiment, the following parameters were chosen a priori:

optical wavelength λ	:	633 nm
fibre diameter $2R$:	100 μm
acoustic frequency f	:	1.5 MHz
larger core V_1	:	2.0
smaller core V_2	:	1.6
maximum intrinsic coupling M	:	0.04

and consequently the following values for the other fibre parameters are required:

numerical aperture NA	:	0.128
larger core radius a_1	:	1.58 μm
smaller core radius a_2	:	1.26 μm
core separation d	:	5.41 μm
cut-off wavelengths of the cores	:	526 nm and 421 nm

For a device length of $L = 30$ cm, an acoustic amplitude of 25.9 nm, corresponding to an acoustic power of 0.6 mW, is required. For $L = 5$ cm, an acoustic power of 21.6 mW would be required. We will see that these are very large power requirements, compared to those available with the null coupler design considered later.

4.4 Effects of variations in the *a priori* parameters

The above design is not the only one which would provide appropriate frequency shifter operation. Other designs are available with different values chosen a priori for the fibre diameter and the core V values, for instance. The process of fibre fabrication is not perfect, and some deviations from any design can be expected. However, by taking measurements at intermediate stages in the fibre fabrication process, deviations in some design parameters which have become fixed (e.g. at the preform stage, the NA , or the ratio of core and cladding diameters) can be compensated for by adjusting other parameters not yet fixed (e.g. the core separation, the ratio of core sizes, or the diameter of the fibre). Fibre design and optimisation does not therefore end with the delivery of a prescription to the fabricator, but continues after each stage of the fabrication process and informs the next stage.

To gauge the influence of the choice of the *a priori* parameters on device performance, the effect of each is considered while the others are kept fixed, with the above values. The values of consequential parameters are adjusted to maintain the *a priori* parameters (e.g. if V_1 is varied, d must be adjusted to maintain M constant). The fixing of V_1 , V_2 and M turns the most complicated functional dependences in (39), (45) and (46) into constants, and so the dependence of derived parameters on λ , f , R and L can be expressed very simply. In particular, the acoustic

power P required to satisfy the maximum conversion condition (44) varies as

$$P \propto \frac{\lambda f R^4}{L^2}. \quad (49)$$

Most strikingly, P varies as the fourth power of R ; hence reduction of the fibre diameter to the minimum consistent with fibre handling is highly desirable. The dependence on λ and f is also notable, indicating more efficient performance at lower acoustic frequencies and shorter optical wavelengths.

The effects of changes in V_1 , V_2 or M cannot be expressed analytically. The dependence of the power P required for maximum conversion is plotted in Fig. 10 against each of V_1 and V_2/V_1 while the other is fixed (note that the ratio of V values is equal to the ratio of core diameters). It is the ratio of V values, rather than the absolute value of either, which is of primary importance here, the main result being that the core sizes should not be too similar. In any case, there is not much scope for variation of the absolute V values, which are effectively restricted between 1.2 and 2.4.

Fig. 11 illustrates the effect of maximum intrinsic coupling M on P , and also on the image sideband suppression. There is clearly a trade-off between power and suppression, with $M = 0.04$ being a good initial compromise.

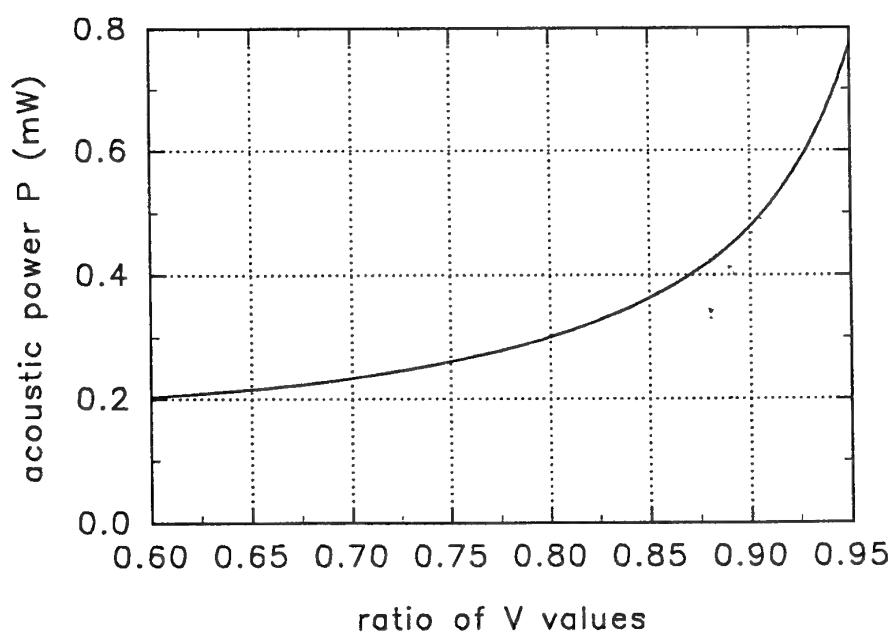
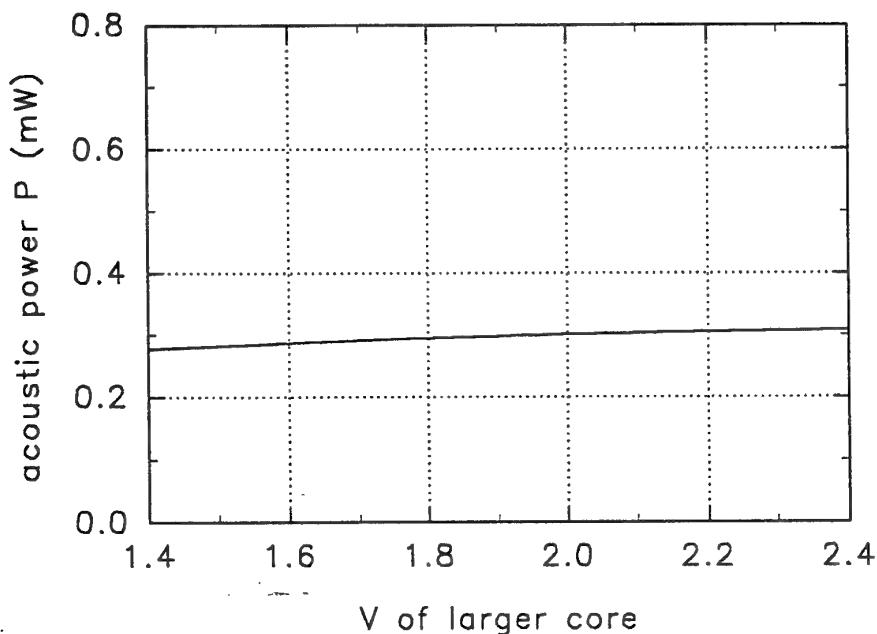


Fig. 10. The variation of acoustic power P for maximum conversion with each of (a) V_b , and (b) V_2/V_b , while the other is kept fixed.

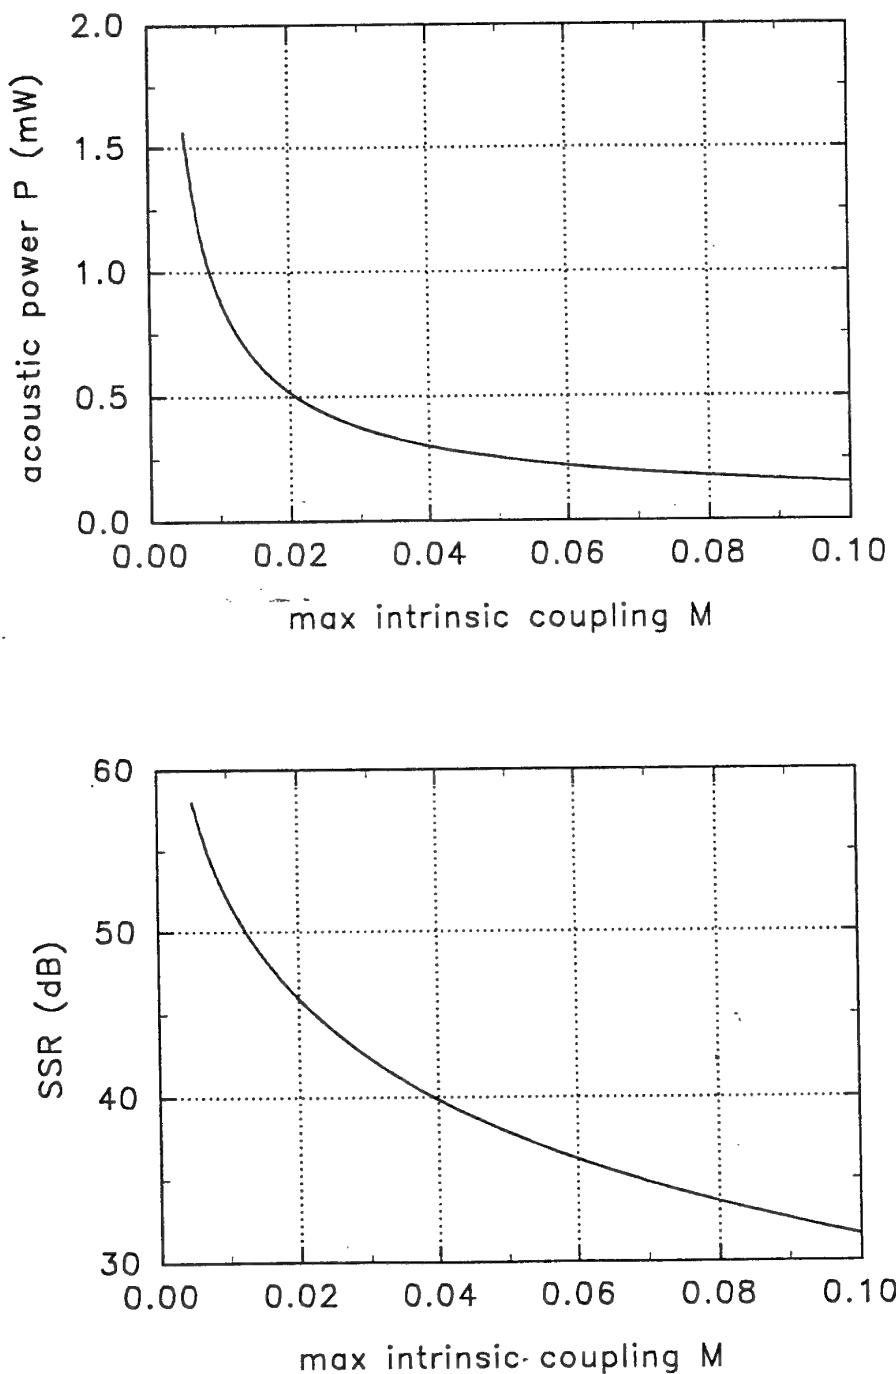


Fig. 11. The variations of (a) acoustic power P for maximum conversion, and (b) image sideband suppression, with the maximum intrinsic coupling M.

5. DUAL-CORE FIBRE: EXPERIMENT

5.1 Fibre fabrication

Dual-core and single-core fibres for the experiment were made at the ORC. The dual-core fibre was fabricated to be close to the design detailed above. A single-core fibre was also fabricated to be splice-compatible with either core of the dual-core fibre.

A pre-existing standard preform was selected and re-characterised using a York P101 preform profiler to determine the *NA* (= 0.113) and the ratio of core and cladding dimensions. These parameters differ from those in the initial design, so the design was modified to allow for this. Three single-core fibres (designated HD015) with diameters of 75, 86 and 97 μm were pulled from this preform. The preform was then cut into two sections, one of which was stretched to reduce the size of its core relative to that in the other, unstretched, section. A flat surface was milled into each preform close to the core, and the two sections were assembled flat to flat to give a dual-core preform (Fig. 12). Three dual-core fibres (designated TC019) were then pulled from this preform (Fig. 13). The fibres were given a range of widths (70, 80 and 90 μm) to allow some flexibility in achieving acousto-optic resonance in the eventual experiment.

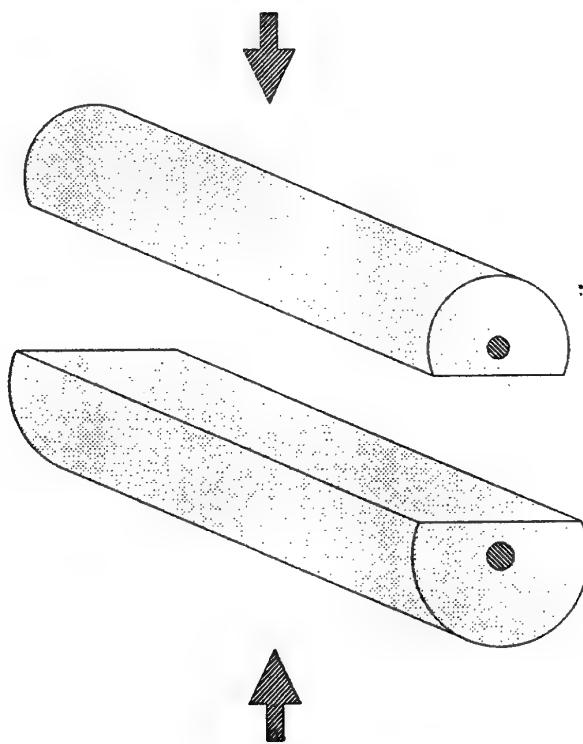


Fig. 12. Assembly of the dual-core preform from milled single-core preforms.

The fibres fabricated as described above should have characteristics close to those of the design. One test for this is a measurement of the cut-off wavelengths of the fibres using the bend-loss

technique. This measurement was made for a piece of single-core fibre with a diameter of 86 μm , and a cut-off wavelength of 585 nm was found. This is markedly different from the expected value of 478 nm. The source of the discrepancy was traced to the P101 preform profiler, which was out of calibration due to the recent removal of ORC labs from one building to another, and which had indicated that the fibre *NA* was lower than it was. The correct *NA* was found to be 0.138, which is consistent with the cut-off measurements. As a consequence, the fibres which had been fabricated from this preform were then predicted to reach acousto-optic resonance at an acoustic frequency of around 2.7 MHz instead of the target of 1.5 MHz. However, since it would be some time before the fibre fabricators would be able to make a second fibre for this project, work proceeded based on this fibre.

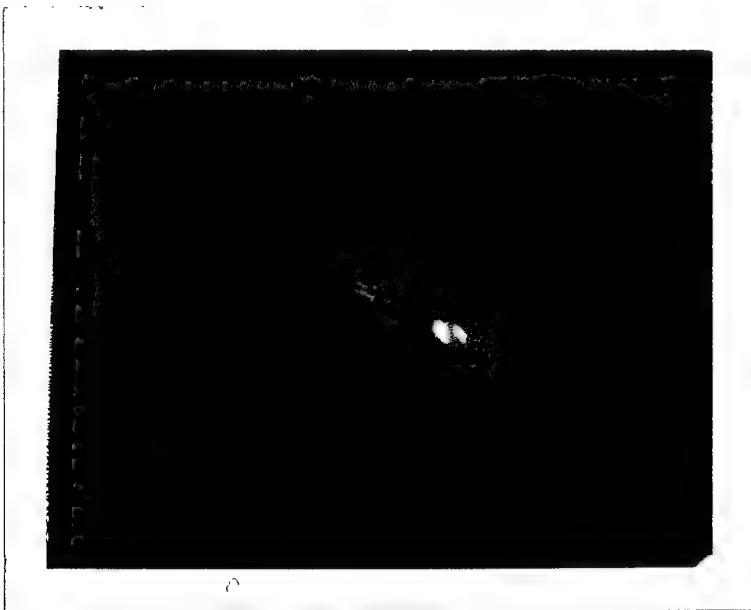


Fig. 13. Cross-section of the dual-core fibre, as seen under a microscope.

A more serious problem with this fibre then came to light. Preliminary inspections of the fibre, such as a test of an output imaging technique to be discussed below, confirmed that light was being guided in both cores. Later, however, at the time of assembly of the experiment, it was found that all remaining samples of the fibre were guiding in one core only. No light was being guided by the second core in even the shortest lengths of fibre. This problem was ascribed to a variation in the depth of milling along the original single-core preform during the construction of the dual-core preform, which apparently penetrated the core along a length of the preform and thus impaired its guiding ability in most of the eventual fibre. It is still something of a mystery how this impairment could be so complete and so sudden. More fibre was made from a small remaining piece of the preform, which was believed not to suffer from this problem. Unfortunately, this was not the case.

A completely new preform was then made, and a fibre pulled from it. This work was further delayed by problems with the newly-installed fibre fabrication equipment. The new fibre was eventually delivered, but work on the dual-core device was then suspended in favour of the more promising taper-based scheme.

The characteristics of the new fibre, and the expected parameters of the acoustic wave necessary to give efficient frequency-shifting at 633 nm, are given below:

fibre diameter $2R$:	$70 \mu\text{m}$
diameter of larger core $2a_1$:	$3.39 \mu\text{m}$
ratio of core diameters a_2/a_1	:	0.83
core separation d	:	$6.7 \mu\text{m}$
numerical aperture NA	:	0.13
maximum intrinsic coupling M	:	0.005
resonant acoustic frequency f	:	0.73 MHz
amplitude for complete conversion u_o	:	87 nm
power for complete conversion P	:	0.9 mW

Particular care was taken during preform assembly to ensure that both cores were not damaged by over-milling. As a result of this caution, the core separation was greater than intended and the maximum intrinsic coupling was expected to be 0.005 instead of 0.04. A greater acoustic amplitude would be needed for maximum conversion, but improved carrier and image sideband suppression would also result.

Simple inspection by microscope of one fibre end, while light from a HeNe laser was launched into the other end, confirmed that all samples of the fibre guided in both cores. A measurement of the fibre's cut-off wavelengths, using the bend-loss technique with cut back, gave cut-off wavelengths of 575 nm and 480 nm (one for each core). The core diameters are in the same ratio.

5.2 Acoustic transducer

The transducer used to transform an RF electrical signal into a longitudinal acoustic wave is a disc made from the piezoelectric ceramic material PZT-4, supplied by Vernitron. This material is designed for high-power operation. The resonant frequency of the disc is determined by its thickness, so a range of discs with resonant frequencies between 1 MHz and 2.5 MHz (i.e. centred around the target frequency of 1.5 MHz) were supplied. Despite the unexpected error in the fabrication of the dual-core fibre, which shifted the predicted frequency of acousto-optic resonance to 2.7 MHz, the 2.5 MHz discs should provide enough acoustic amplitude to drive the frequency shifter.

The acoustic wave from the PZT disc is concentrated by a tapering glass horn before being imposed on the fibre. The horns used are a by-product of the fibre fabrication process; a horn-shaped structure remains on the end of a preform after a fibre has been pulled from it. (An attempt was made to make horns in the optics lab by directing an oxy-butane torch onto a weighted glass rod, but it was not possible to form a horn which is not excessively buckled, even if the burner is rotated about the rod.) These horns are made from a lead silicate glass instead of silica as first proposed. This is because the lead silicate glass has a much lower softening point, which greatly helps the process of fusion between the tip of the horn and the fibre. The wide end of the horn is cut from the host rod using a glass saw and then polished flat. The narrow end is cleaved flat using an alumina blade to score the glass before inducing a clean brittle fracture.

Ideally, all the acoustic energy generated in the PZT disc is transmitted to the horn, and none is passed to the support onto which the disc is fixed. Little loss of energy into the support will occur if it is made from a material with an acoustic impedance which is grossly different to that of the disc. Air would provide the greatest mismatch, but air of course does not provide support. An appropriate material is a block of silicone rubber which, with an acoustic impedance of 1.07×10^{-6} kg m⁻² s⁻¹ compared to 34.5×10^{-6} kg m⁻² s⁻¹ for PZT-4, would reflect 88 % of the incident acoustic power. A better technique, however, is to mount the disc on an aluminium support which is half an acoustic wavelength thick. This gives an effective acoustic impedance of zero as seen by the disc. In addition, the metal acts to some extent as a heat sink for the disc, allowing it to be driven harder before heat-induced depoling of the material occurs. This is the technique adopted. It has the disadvantage of restricting the range of acoustic frequencies which may be used to those for which the aluminium plate is a half wavelength thick. For 2.7 MHz, the half wave thickness of aluminium is 1.19 mm. An aluminium plate a few millimetres thick is taken, and a 10 mm diameter blind hole was bored to a distance of 1.2 mm from the back surface. A small hole was drilled through the thinned region. The PZT disc was fixed to the surface of the thinned region with epoxy. Disc and plate were pushed together while the epoxy cures, to minimise the thickness of the epoxy layer and improve thermal contact. The small hole accepts the protruding solder joint of the electrical wire supplying the electrode on the bottom of the disc (Fig. 14).

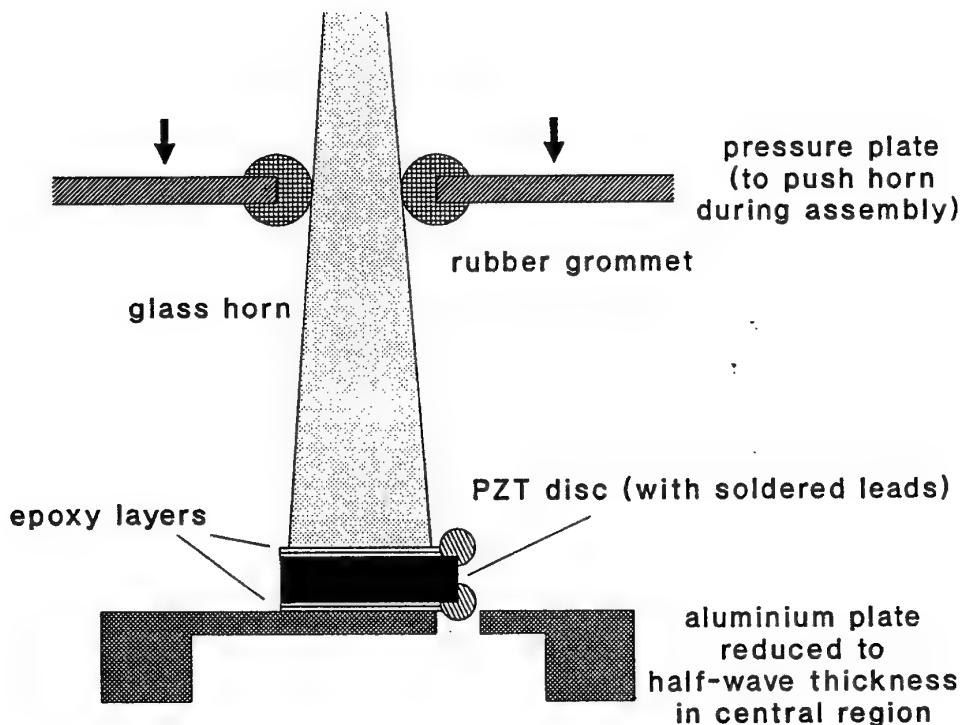


Fig. 14. The construction of the transducer assembly. The pressure plate is removed after the horn has been fixed to the disc.

The horn was also attached to the upper surface of the disc with a layer of epoxy cured under pressure. It is particularly important to minimise the epoxy thickness here, since it is a potential source of acoustic reflection. (An attempt was made to attach the horn using salol, a crystalline material with a low melting point which is often used in temporary acoustic assemblies, but the

joint formed was too weak. If necessary, disc and horn can be joined instead using indium solder, but the PZT material would then have to be re-poled.) To force the horn onto the disc it was necessary to apply pressure via a metal plate with a rubber grommet, to push the horn without damaging it. PZT disc area in excess of the footprint of the horn is wasteful of supplied RF power, so surplus segments of the disc were sawn off prior to attachment of the horn (slowly, to avoid too much frictional heating) using a diamond saw.

The transducer was driven by an RF signal generator, with an amplifier capable of supplying up to 5 W electrical power if necessary. This supply was matched as usual to 50Ω impedance. However, the electrical characteristics of the transducer were very frequency dependent and deviated greatly from a 50Ω resistive impedance (the impedances were measured using a RF impedance matching network in the Department of Electronics and Computer Science). An LC matching circuit was therefore built to increase the efficiency of coupling of electrical power into the transducer. For most of the frequency range of interest, the transducer assembly (including support plate and horn, but excluding the fibre) had a capacitative reactance, so a series inductor was introduced to cancel this out. To permit some tuning of the matching reactance (to compensate for the strong frequency dependence of the transducer's impedance), excess inductance was introduced-together with a variable capacitor also in series. To match the residual resistive component of the impedance, a small transformer was wound. In fact the reactive elements just described were placed within the primary circuit of the transformer; this still canceled the reactive impedance while allowing components with smaller values to be used (Fig. 15). Only approximate values for these components could be determined initially, because the exact electrical frequency with which to drive the transducer is unknown until the entire frequency shifter is assembled and the frequency of acousto-optic resonance found directly.

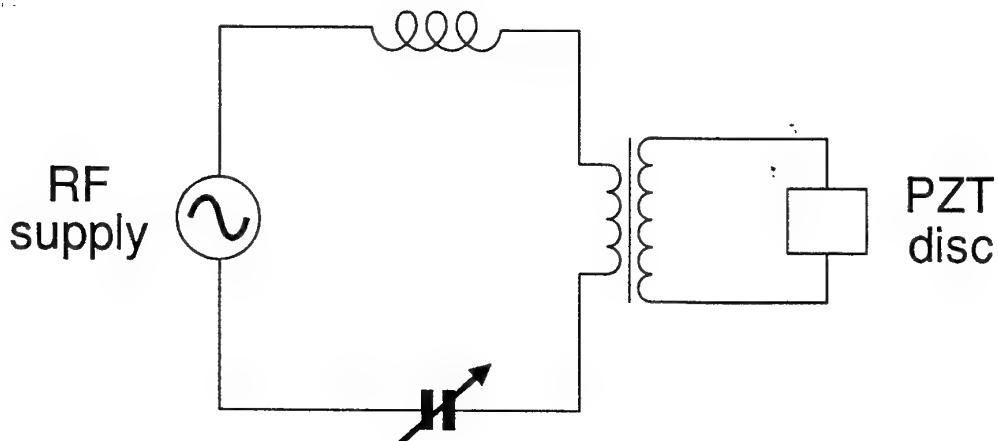


Fig. 15. The electrical impedance-matching circuit.

5.3 Fibre splicing

The fibre holders consisted of two 4 cm lengths of glass capillary, held horizontally, in line, and about 40 cm apart, in aluminium blocks. The fibre was passed through the bores of the capillaries. A 30 cm length of the dual-core fibre, and a short length of the single-core fibre, were stripped of their coatings after immersion for a minute or so in a stripper solution (50:50 acetone and dichloromethane), and cleaned. The fibre ends were cleaved. The fibres were then

spliced together using a standard single-mode fibre fusion splicer.

A standard splicer does not include facilities for splicing dual-core fibre (if indeed any splicer does). As well as the familiar requirements that the splice be low-loss and robust, we have the extra requirement that light from the single-core fibre is coupled efficiently into only one core of the dual-core fibre. In this experiment it is particularly important that none of the light is launched into the second core. Unfortunately, the dual-core fibre is non-circular and neither core is central, so the action of surface tension in the liquified fibre joint as splicing takes place (which tries to centralise and circularise the fibres) ruins the desired alignment if allowed to proceed as normal. In view of these difficulties, a priority is put on the achievement of low crosstalk of light into the second core, rather than a low insertion loss overall.

The cleaved fibres were loaded in the splicer as usual. Light from a HeNe laser was launched, via a set of neutral density filters and/or a grossly inefficient launch (for eyesight protection), into the single-core fibre. The other end of the dual-core fibre was mounted in a fibre end-inspection microscope, so that it was possible to see when the fibres were aligned and particularly when minimum power leaked into the second core. After alignment, the fusion arc was operated in order to join the fibres, but to avoid the detrimental effects of surface tension on the alignment of the fibres, the mains power to the splicer was switched off part-way through the arc's action. In this way, reasonably strong splices with low crosstalk to the second core could be made.

The insertion loss of the splice, and hence of the eventual frequency shifter device, could be measured using a cut-back technique whenever necessary, comparing the total output power from the device with the light emerging from a piece of the single-core fibre cut from the input.

5.4 Fibre rotational alignment

The spliced fibres were carefully removed from the splicer and suspended between the fibre holders with a minimum of slack and with the bare portion centrally located. Hence the splice lies towards one end. Before the horn is fused to the fibres, it is first necessary to rotate the fibre so that the cores in the dual-core fibre lie vertically, one directly above the other. Here we can take advantage of the non-circularity of the dual-core fibre, which makes it possible to determine the rotational positioning of the fibre by side inspection. By gaining familiarity with the appearance of a stub of the fibre under different rotations, verifying the exact rotation of the fibre by looking at its end, it was possible to align the spliced fibre mounted for the experiment by side inspection only, with the aid of a down-directed table lamp and a 7 \times hand magnifier based on the table. In fact, because this fibre was so non-circular, it is possible that the state of rotation of the fibre in regions away from the horn and splice is not critical; the fibre may be sufficiently acoustically-birefringent for the plane of the flexural wave to rotate with the fibre. Other schemes for aligning the fibre were identified, should future fibre samples have proven to be more featureless externally. For example, the diffraction pattern of the fibre in a transversely-directed laser beam depends on the fibre's orientation.

5.5 Fixing of fibre and horn

As discussed in an earlier section, the horn is fixed to the single-core fibre near the splice, to ensure that the flexural wave generated is unidirectional along the dual-core fibre. It is necessary to provide a join which is strong and acoustically efficient, yet which does not distort

6. NULL-COUPLER FREQUENCY SHIFTER

6.1 Description

While we were waiting for the new dual-core fibre to be made and characterised, the alternative possibility of a frequency shifter at the narrow waist of a fused tapered fibre coupler was explored.

In this context, a taper is a bare fibre, a part of which has been stretched and narrowed in a flame. To form a fused tapered coupler, two bared fibres are stretched and fused together to form a common narrow "waist" structure, which is linked to untreated portions of fibre by taper transitions (Fig. 16). Optically, this waist is a single composite waveguide, supporting its own set of normal modes. The fibre cores are too small to guide light, and an incoming light wave spreads into the cladding, to be guided by total internal reflection at the cladding-air interface.

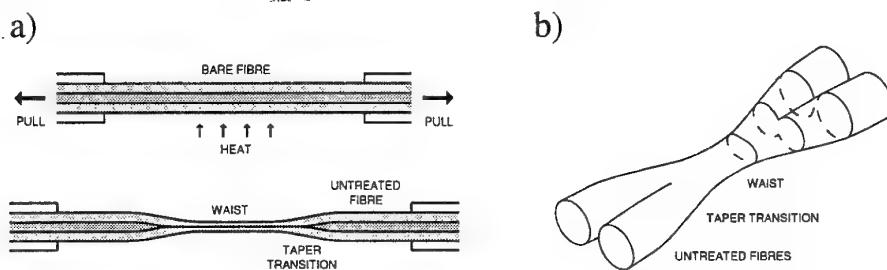


Fig. 16. Schematic diagrams of (a) the formation of a fibre taper, (b) a fused tapered coupler.

For conventional beam-splitting applications, a coupler is usually made from two identical fibres. Light entering one fibre excites two modes of the composite waveguide at the coupler waist, which interfere to cause coupling. Here we consider a coupler made from two markedly dissimilar fibres, so that light entering in one fibre core excites only one normal mode in the waist. If the taper transition is gradual enough, light in the fibre with the greater propagation constant evolves into an even mode (the fundamental mode of the waist), while light in the other fibre evolves into an odd mode (the second mode of the waist). At the downstream end of the coupler, the light returns to the fibre core from which it emerged. The passive device does not actually function as a beam-splitter, so we call it a "null" coupler. As a passive component, a null coupler is useless. It is equivalent in passive function to a pair of unrelated fibres held together in parallel by sticky tape! Consequently, null fibre couplers have received little attention, and there is no known report of one ever having been made by the taper technique (though their behaviour has been described for planar waveguides⁷). However, unlike a pair of unrelated fibres, in a null coupler the light waves from the two input fibres interpenetrate and overlap in the circularly-fused taper waist. If a flexural acoustic wave is imposed on the coupler with a wavelength matching the beatlength of the two normal modes of interest, resonant coupling can take place between them, thus permitting frequency-shifted light to enter the other fibre.

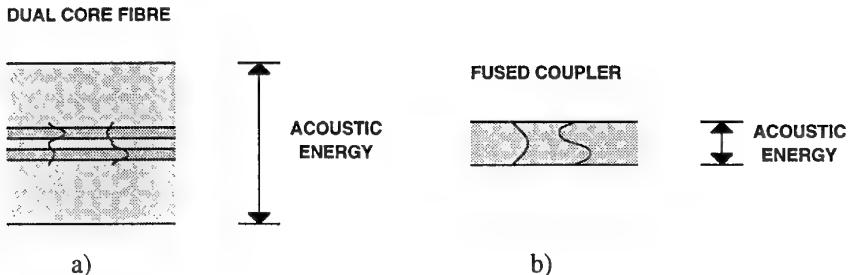


Fig. 17. The acousto-optic overlap in (a) a dual core fibre, where the acoustic energy at the edge of the fibre is wasted, and (b) a coupler waist, where all of the acoustic wave interacts with the light. Devices based on two-mode fibre also suffer from the same disadvantage as (a).

The acousto-optic interaction is entirely analogous to that in a dual-core fibre, or indeed in a two-mode fibre. However, in contrast with such fibres, here the normal modes fill and share all the available glass volume, despite each being excited from one input fibre only (Fig. 17). The whole of the acoustic wave interacts with the light, and so much less RF power is needed to drive the device. This crucial difference arises because the coupler, and hence its modal properties, varies longitudinally. This permits a gradual transition between an optimal acousto-optic interaction region at the coupler waist, and standard single-mode fibres.

The waist of an ordinary beam-splitter coupler is typically 20-30 μm wide and about 10 mm long, with each taper transition also about 10 mm long. Despite its small size, such a coupler is surprisingly robust, provided its surface is not damaged by handling; they have been used as fibre beam splitters for many years now. A coupler for our frequency shifter experiment would be even thinner and longer, but is still not expected to suffer from excessive mechanical weakness.

6.2 Advantages over dual-core fibre

Such a frequency shifter has a number of major advantages over a device based on dual-core fibre (or indeed any fibre where light is guided in a buried core):

- The two interacting optical modes fill the same volume of fibre as each other and as the whole of the acoustic wave, so a strong acousto-optic interaction is expected. This means that the power requirements for the transducer which drives the acoustic wave will be very small.
- The excess of acoustic power now available permits a much reduced acousto-optic interaction length. This gives a more compact device with possibly a greater bandwidth.
- The taper transition functions as an acoustic concentrator, further reducing drive power requirements. Much less acoustic power is needed for a given displacement in a narrow coupler waist than in a full-size fibre.
- In principle, the coupler can be made with near-zero intrinsic coupling, without affecting the acoustic power requirements. This makes possible an optical output with near-zero power at the carrier and image-sideband frequencies. In the dual-core fibre, there is a compromise between power requirement and carrier suppression.
- The device automatically has four fibre ports. No splices or mode-convertors are necessary.
- Because the coupler waist can be made very much narrower than a normal handlable

fibre, much higher operating frequencies are feasible.

- No special fibres need to be made. The exact nature of the fibres is unimportant, provided they are dissimilar enough and are both single-mode at the operating wavelength (e.g. 633 nm). However, a general requirement for making low loss couplers is that the fibres have a "matched" cladding, with no refractive index depressions.
- The interaction region is defined by the coupler waist, the dimensions of which can be easily controlled when the coupler is made. This gives considerable freedom in device design, which is not constrained by the particular single-mode fibre which may happen to be available. A given pair of standard single-mode fibres can be used to provide devices for any of a wide range of RF frequencies and optical wavelengths; the particular frequencies and wavelengths of operation are only fixed when the coupler is made.
- If the acoustic wave is excited at one of the untreated fibres, it is unidirectional in the interaction region. This eliminates one source of image-sideband frequency-shifted light in the output.
- Device bandwidth can be increased by making the coupler waist slightly non-uniform.
- The resulting frequency shifter would supersede existing technologies, being entirely novel.

6.3 Making tapers and couplers

A coupler fabrication system (the "coupler rig") has been built. It is based on a design used successfully by Tim Birks and his former colleagues at the University of Limerick. The two stripped fibres are held in V-grooves mounted on two carriages, the elongation stages, sharing a common rail (Figs. 18 and 19). The elongation stages can be moved apart at a rate of typically 3 mm/min by geared stepper motors. This serves to stretch the fibres. Simultaneously, the fibres are softened and fused by a small oxy-butane flame. The burner is mounted on a carriage on a parallel rail and is "brushed" to and fro along the fibres. The speed of travel is constant and is typically 400 mm/min, i.e. fast compared with the elongation stages. Thus a length of the fibres equal to the burner's travel is heated uniformly and will be uniformly reduced in size when stretched. The heating and stretching processes can be stopped when the appropriate coupler length and width are obtained. The progress of coupler formation can be monitored *in situ* by launching light into one end of a fibre and detecting the light emerging from the other ends of both fibres.

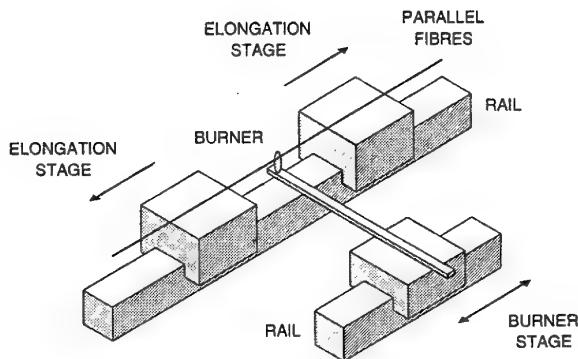


Fig. 18. A schematic diagram of the coupler rig.

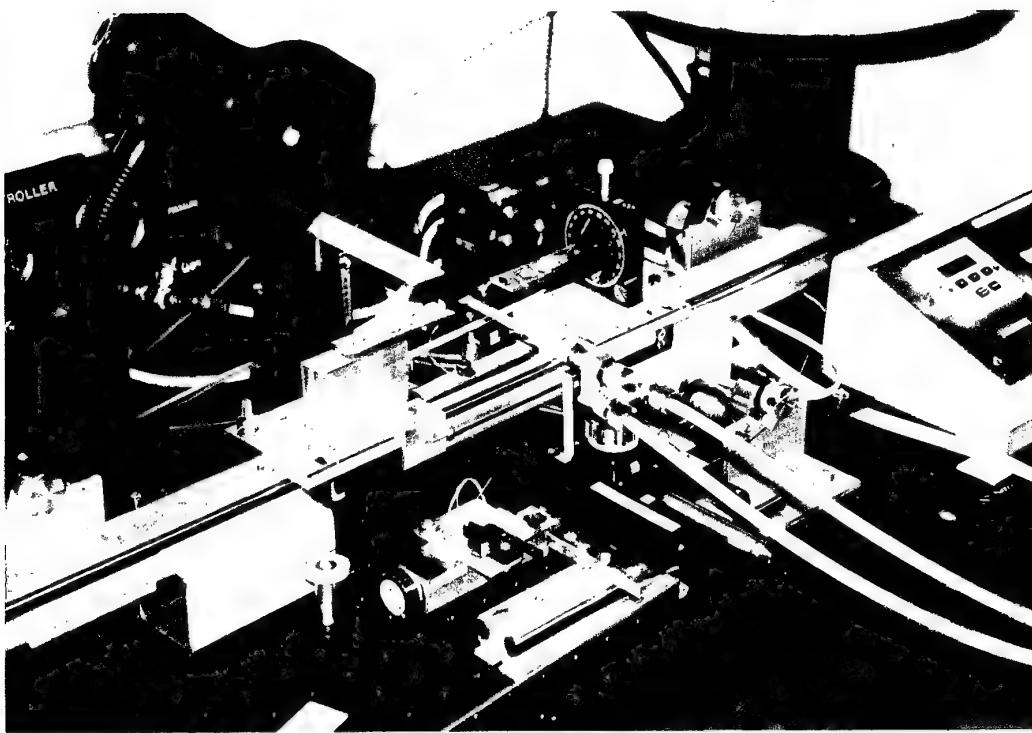


Fig. 19. A photograph of the coupler rig.

The same rig can also be used to make single tapers. In this case, just one fibre is held in the V-grooves.

As a coupler is formed, the local diameter of the fibres changes, as does the degree of fusion (the extent to which the two fibres have melted together). The final coupler therefore exhibits transitions in cross-sectional size and shape.

The control of the diameter of a single taper is now well understood (for a coupler, "diameter" must be interpreted as the diameter the coupler would have if both fibres were completely fused together to give a structure with a circular cross-section). A satisfactory theoretical model has been developed, and is found to be effective⁸. The taper's waist and transitions must be considered separately. For each pass of the burner along the fibre(s), every part of the taper waist is heated and stretched identically. The waist therefore has a uniform and repeatable diameter. This is important for an acousto-optic device, because the condition for acousto-optic resonance depends critically on the size and shape of the waveguide. The taper transitions are formed by elements at the ends of the waist being pulled out of the heated region. By changing the burner's distance of travel during taper elongation, the length of each element can be adjusted, thus permitting complete control of the longitudinal diameter profile of the taper. The burner travel is of course easily controlled. Less controllable parameters, such as gas mixture and flow rate, or burner geometry and height, have little effect on the longitudinal diameter profile.

The cross-sectional shape of the coupler, and particularly its longitudinal variation, is not satisfactorily modelled. It is also difficult to control experimentally since, unlike the local taper

diameter, it does depend on gas and burner parameters, which can vary to some extent not only from one rig to another, but also from day to day. Cross-sectional shape is difficult even to measure; it is necessary to cleave the coupler and examine the end under a microscope, a process which of course destroys the coupler, and in any case only reveals the cross-section at one point. However, there is one method by which a known and repeatable cross-sectional shape at the coupler waist can be achieved. As coupler elongation proceeds, the fibres at the waist become progressively more closely fused under the influence of surface tension. The shape changes from two touching circles, through a "dumb-bell" shape, to elliptical and finally circular shapes, given enough time and heat. Once the cross-section is circular, it will remain so, since this shape gives minimum surface energy for a given cross-sectional area. Hence, provided that more than enough time and heat have been supplied, we can know that the cross-section at the waist is circular despite any small variations in fabrication parameters. The coupler should therefore be stretched at a low speed in a hot flame. Appropriate conditions can be found by making a series of test couplers and cleaving them to confirm that their cross-sections are circular. We can then be confident that further couplers produced under similar conditions will also be circularly fused. The circular cross-section has one additional advantage: it is the simplest to model theoretically for its optical and acoustic properties, which greatly simplifies the design procedure.

The detailed longitudinal variation of the cross-section in the taper transitions (from narrow and circular at the waist to two separate fibres at the untreated ends) is still not predictable, but this is not critical. Provided the transitions are such that the coupler functions as a null coupler, their shape is unimportant.

6.4 How a taper works

It is usually understood that a single-mode fibre can support just one mode of light (or two if the two independently-polarised states are considered distinct). This fundamental mode is guided by total internal reflection at the core-cladding boundary. It is assumed that the fibre cladding extends to infinity, which is a good approximation when considering the properties of the fundamental mode. However, a real fibre has a finite cladding. If the medium surrounding the fibre has a low index, like air, then a large number of additional modes are also supported. These are guided by total internal reflection at the cladding-air boundary, and are called cladding modes. Cladding modes are suppressed in a standard fibre by a lossy high-index polymer coating around the fibre, which absorbs, or "strips", cladding mode light.

A single-mode fibre taper is a longitudinally-varying waveguide, and is uncoated. As a light wave in the fundamental mode of the untreated fibre enters the taper transition, it sees the fibre core reducing in size. As the core reduces in comparison with the wavelength of the light, its effect on light guidance decreases. Consequently, the local fundamental mode spreads out relative to the core and extends further into the cladding. If this process is allowed to continue, the mode spreads as far as the cladding-air boundary, which becomes more important to light guidance than the diminishing core. Eventually, the fundamental mode is guided by the cladding-air boundary, with the core playing little role; it has become a cladding mode. In most fibre tapers, this takes place by the time the diameter of the waist has been reduced by a factor of four or so. At the second taper transition, the process is reversed; the fundamental mode is recaptured by the enlarging core. This process is depicted in Fig. 20.

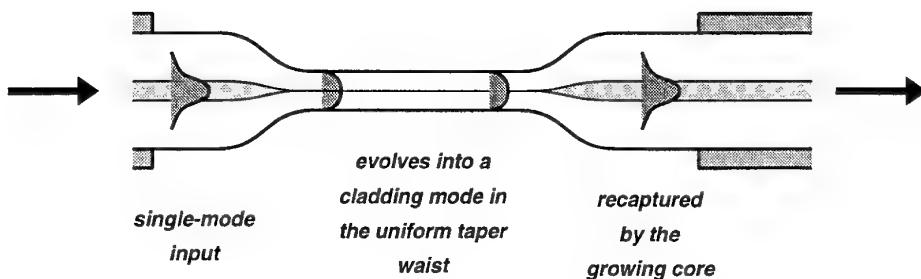


Fig. 20. Evolution of the fundamental mode through a tapered fibre.

The waist also guides many higher-order cladding modes. If the taper transition is abrupt, the light wave cannot follow the fundamental mode as it spreads, so there is mode-coupling to the higher-order cladding modes. Any light remaining in these cladding modes downstream of the taper is not recaptured by the core. It remains in the cladding until the fibre coating resumes, and then is stripped by the coating. Hence abrupt taper transitions generally cause loss of light. However, if the transition is sufficiently gradual, there is negligible mode coupling; the light in a given mode remains in that mode throughout the taper, even if the field distribution of the mode changes. Such a transition is called "adiabatic". Theoretical criteria specifying whether a given transition is adiabatic have been extensively described in the literature⁹. A taper in a normal matched-cladding fibre is usually adiabatic if the hot-zone is longer than about 6 mm, so low-loss tapers are fairly easy to make.

A primitive acousto-optic device can be made from a single taper. In the waist of an adiabatic taper in a single-mode fibre, only the fundamental mode is excited. If the taper is passive, the light goes straight through the taper and out the other side, little changed. However, a flexural acoustic wave with the right frequency and amplitude will couple and frequency-shift light from the fundamental mode to the second mode at the uniform taper waist only. The acousto-optic interaction is the same here as for the dual-core device, or the two-mode fibre device developed at Stanford. However, because the two modes occupy the same space as each other and as the acoustic energy, coupling is very efficient. The taper transition also serves as a concentrator for an acoustic wave generated by a transducer attached to the un-narrowed fibre. The light in the second mode emerges from the far side of the taper, remains in the cladding, and is stripped off by the fibre coating where it resumes (Fig. 21). Thus the frequency-shifted light is in fact lost, and little or no light emerges from the fibre end downstream of the taper. The device is therefore best described as a switch, filter or modulator, rather than as a frequency shifter.

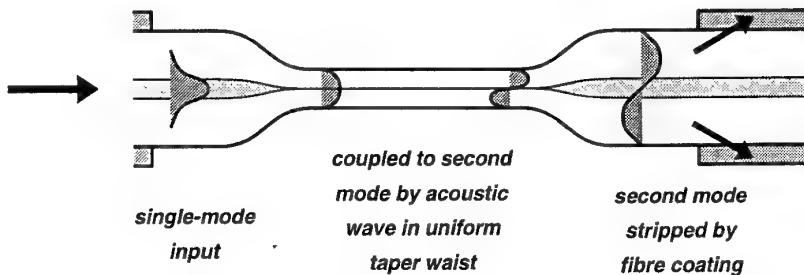


Fig. 21. The acousto-optic interaction in a tapered fibre. The acoustic wave is assumed to have the right frequency and amplitude for complete conversion on resonance.

The frequency-shifted light could be recovered if the fibre was two-moded; the light coupled into the second mode would be recaptured by the fibre core downstream of the taper, in the same way as the fundamental mode. The output of the device would be the same as for the Stanford two-mode fibre frequency shifter: a shifted wave in the second mode, together with a residual wave at the original ("carrier") frequency in the fundamental mode. It would then be necessary to extract the power in the second mode and convert it to a fundamental mode, for compatibility with single-mode fibre. Imperfections in this process lead to contamination of the output wave by light at the carrier frequency. An alternative means to extract the frequency-shifted wave is to use a null coupler.

6.5 How a coupler works

A weakly-fused coupler, with a cross-section resembling two touching circles, can be regarded as being an interacting pair of single tapers. Since the light in a taper waist is guided at the outer boundary, it can leak from one taper into the other, thus splitting the light between the two fibres. However, if the cross-section is strongly fused, the two individual fibres lose their separate identities, and the coupler waist must be treated as a single composite waveguide supporting its own set of normal modes. If each separate taper transition is gradual enough to be adiabatic on its own, then an incoming light wave in the fundamental mode of *one fibre* will excite some combination of the fundamental and second modes of *the coupler waist* (Fig. 22). An incoming wave in the other fibre will excite the orthogonal combination of the modes of the waist.

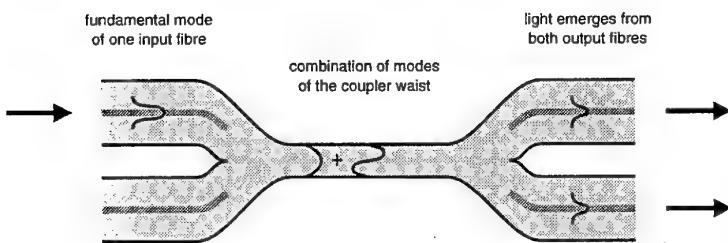


Fig. 22. The excitation of modes in a standard fused tapered coupler.

The fundamental and second modes of the coupler waist have different propagation constants, so as they propagate along the waist they acquire a phase difference. Unless the length of the coupler is such that the phase difference is a whole number of complete cycles, there will be some light at the far end of the coupler in the combination of modes corresponding to light in the second fibre. Thus some light is coupled into the second fibre.

Although the coupler is *weakly* adiabatic in the sense that there is no mode-coupling in either individual taper, in general the coupler as a whole is not *strictly* adiabatic, because an input in one mode excites more than one mode within the coupler waist. Consequently, there is more than one output mode: light emerges from both fibres.

In the absence of loss, the fraction P_c of light coupled into the second fibre can be expressed as

$$P_c = M \sin^2(\Phi) , \quad (50)$$

from (7), the power remaining in the first fibre being $P_t = 1 - P_c$ by power conservation. Thus P_c depends on two factors. Firstly, it depends on Φ , which gives the particular phase difference

acquired by the two modes as they propagate along the coupler waist. This varies with changes in the length of the waist, its width, the refractive index of the fibres (which depends on temperature), the mechanical state of the coupler (which could be under extensional, bending or torsional strain), and the wavelength of the input light. As Φ is varied by tuning any of these parameters, the coupled power will vary between zero and M .

Secondly, P_c depends on the maximum intrinsic coupling M , which is determined by the ratio ξ of powers excited in the fundamental and second modes of the coupler waist, (8). If the fibres comprising the coupler are identical, the two modes of the coupler waist are excited equally and $\xi = 1$. As a result, $M = 1$ always, and it is possible for all the light to be coupled into the second fibre, given a suitable value of Φ . If the fibres are dissimilar, M can have any value between 0 and 1. It is decreased by increasing dissimilarity, decreasing degree of fusion, and more gradual transitions. In the limit of extreme dissimilarity, $\xi \rightarrow 0$ or ∞ , with only one mode being excited in the waist by an incoming wave in one fibre only. Such a coupler has $M = 0$ and does not couple any light for any Φ ; it is a null coupler. A null coupler is strictly adiabatic, because only one mode is excited in the coupler waist when there is only one input mode. Note that a null coupler must have $M = 0$; although a coupler with $M \neq 0$ and Φ an integer multiple of π also exhibits zero coupling, it is not a null coupler.

For our acousto-optic device, we need an input wave in a single-mode fibre to excite just one mode of the interaction region, which is the circularly-fused coupler waist. Clearly, it is the null coupler which will do this. In the absence of an acoustic wave, light entering one fibre will travel through the coupler in only one of its modes, and will emerge unchanged from the other end of the same fibre (Fig. 23a). If the coupler supports a flexural acoustic wave of the right amplitude and frequency, there is resonant coupling between this mode and the other, with a frequency shift. The mechanism of coupling is identical to that in the primitive acousto-optic device based on a single taper. However, this time the new mode corresponds to light in the other fibre; the frequency-shifted light is neither lost nor entangled with unshifted light, but emerges from the other output port of the coupler (Fig. 23b). Any residual unshifted output light emerges from the first fibre, and the device is automatically a true four-port device.

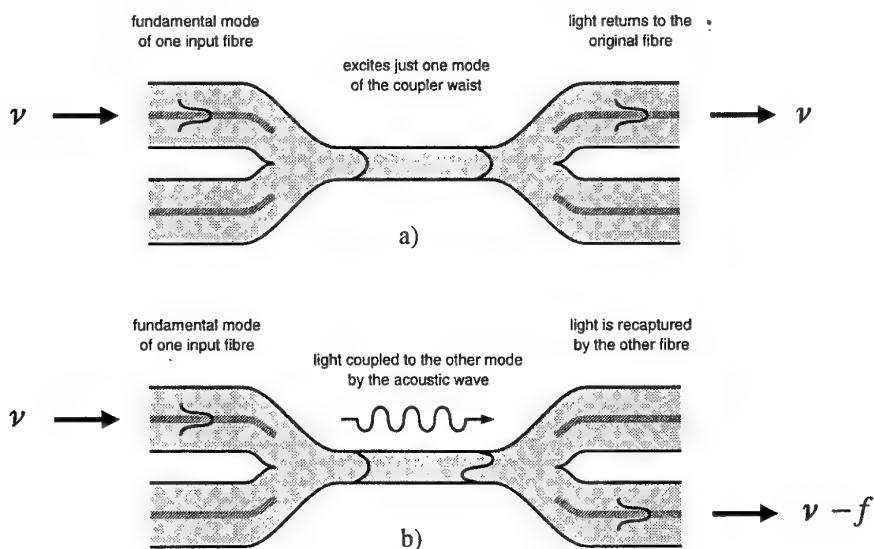


Fig. 23. (a) The evolution of the fundamental mode of one fibre through a null coupler. (b) The acousto-optic interaction in a null coupler.

7. NULL COUPLER: THEORY

7.1 The taper transitions

An exact analysis of the behaviour of a fused tapered coupler is not possible, because of the complexity and longitudinal variation of the waveguide. In particular, the optical behaviour of the taper transitions in the coupler is not only impossible to model satisfactorily, but the exact variation of the coupler's cross-sectional shape along the transitions is in general unknown. However, the detailed optical properties of the transitions can be ignored provided three conditions are satisfied:

- the transitions provide a low-loss optical connection between the untreated fibre(s) and the taper waist (weak adiabaticity);
- the transitions are such that the coupler approximates a null coupler (strict adiabaticity);
- the waist has a circular cross-section.

The transitions then serve simply as low-loss means to transfer light between an untreated fibre and predominantly one mode of a circular taper waist. If this is demonstrated experimentally for a given coupler, it is not necessary to analyse the transitions theoretically. The maximum intrinsic coupling M is thus treated as an experimentally determined quantity.

The first two conditions can be tested by observation of the coupler's optical behaviour. Low loss and null splitting must be exhibited for more than one set of parameters; even if there is mode coupling, it is possible for there to be a whole number of mode-coupling beats along the coupler ($\Phi = N\pi$ in (50)), giving apparently low-loss and null behaviour. A range in parameters can be furnished by varying the wavelength of the input light, deforming the coupler mechanically by bending, stretching or twisting, or simply by observing the coupler's response as it is lengthened during its fabrication.

Circularity of the waist can only be proved destructively by cleaving it. However, once test couplers have been made with confirmed circular waists, we can be confident that the next coupler made under the same conditions will also have a circular waist.

One important result follows from the assumption of strict adiabaticity. This concerns which of the modes of the taper waist is excited by a light wave in a given input fibre, which in turn determines the sign of any frequency shift. The fundamental mode is always that with the greatest propagation constant. If the whole coupler is strictly adiabatic, then the fundamental mode of the *pair of* untreated fibres must evolve into the fundamental mode of the waist, and the second mode of the *pair of* untreated fibres must evolve into the second mode of the waist. There are two guided modes of the *pair of* fibres, namely the two fundamental modes of the *individual* fibres. That with the greater propagation constant is the fundamental mode of the *pair of* fibres, while that with the lesser propagation constant is the second mode of the *pair of* fibres. Hence a light wave in the fibre with the greater cut-off wavelength excites the fundamental mode of the coupler waist, while a wave in the other fibre excites the second mode of the waist.

Image sideband and carrier suppression in the second output fibre depend as before on the maximum intrinsic coupling M . This can in principle be zero (perfectly null) if the constituent fibres are dissimilar enough and the taper transitions are sufficiently gradual. Coupling to the image sideband can also be caused by the presence of a backward-travelling acoustic wave, generated by reflection of the forward acoustic wave at the far end of the coupler.

7.2 The taper waist

The important optical properties of the taper waist can be well modelled by neglecting the residual cores of the constituent fibres, and assuming the cross-sectional shape of the fused fibre structure to be circular. Also, the waist is highly multi-mode (the two modes of interest to us being the lowest-order modes), so a scalar-wave (LP mode) approximation can be used and the outer boundary can be treated as being field-excluding, like an ideal conductor. The analysis of the acousto-optic interaction in the waist of a null coupler is identical to that for a single taper, since the waveguide in the interaction region is the same in both cases, namely a cylinder of cladding glass surrounded by air.

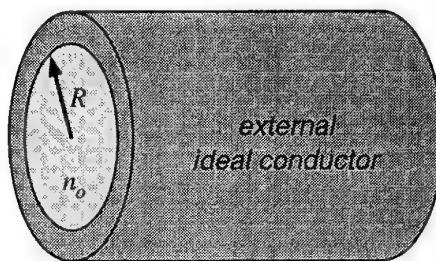


Fig. 24. Model for the waist of a taper or coupler.

The coupler waist is modelled by a cylindrical dielectric rod with a radius R and a uniform internal refractive index n_o , surrounded by an ideal conducting medium (Fig. 24). This is a good approximation if the coupler waist is highly multi-mode. The fundamental (even) and second (odd) modes are readily shown to have the normalised field distributions

$$\psi_e = \frac{1}{R\sqrt{\pi}} \frac{J_0(j_{01}r/R)}{J_1(j_{01})}$$

and

$$\psi_o = -\frac{\sqrt{2}}{R\sqrt{\pi}} \frac{J_1(j_{11}r/R) \cos \varphi}{J_0(j_{11})}, \quad (51)$$

and their beatlength is

$$L_B = \frac{8\pi^2}{(j_{11}^2 - j_{01}^2)} \frac{R^2 n_0}{\lambda}, \quad (52)$$

where r and φ are polar co-ordinates, and J_0 and J_1 are Bessel functions with first roots $j_{01} = 2.405$ and $j_{11} = 3.832$ respectively.

The resonance condition (21) must also be satisfied, with L_B from (52) and Λ from (9); in the thin coupler waist, the low frequency limit of the acoustic dispersion relation is guaranteed valid. The resonance condition is then expressed as a relation between R , λ and f :

$$\left(\frac{\Lambda}{L_B}\right)^2 = \frac{(j_{11}^2 - j_{01}^2)^2}{64\pi^3} \frac{c_{ext} \lambda^2}{n_0^2 R^3 f} \equiv 1 . \quad (53)$$

This simple equation is of fundamental importance in the design of taper-based acousto-optic devices. It tells us the waist radius R needed for the device to give a resonance for light of wavelength λ with an RF drive frequency of f (and hence a frequency shift of f).

The acoustic amplitude u_o for maximum conversion is given by (22) in terms of the overlap integral (23). Using (51), we get

$$I_O = \frac{2\sqrt{2} j_{01} j_{11}}{(j_{11}^2 - j_{01}^2)^2} R , \quad (54)$$

and hence

$$u_o = \frac{(j_{11}^2 - j_{01}^2)^2}{16\sqrt{2} \pi j_{01} j_{11}} \frac{c_{ext} \lambda}{n_o (1 + \chi) f L} . \quad (55)$$

The corresponding power follows from (14). Unlike (44) for the dual-core fibre, (55) has no M dependence, because M in null couplers is a property of the taper transitions, not the interaction region at the waist.

As an example calculation, consider a coupler made from silica fibres carrying light from a HeNe laser at 633 nm. For acousto-optic resonance at $f = 1.5$ MHz, the waist radius required by (53) is $R = 3.1 \mu\text{m}$, which is a small but possible value for a fused tapered coupler. For a waist length of $L = 50$ mm, the required acoustic amplitude is $u_o = 5.2$ nm and the required acoustic power is just $P = 45$ nW.

The above simple analysis can be elaborated by considering an outer medium of air, the presence of residual fibre cores, or a non-circular cross-section. The results obtained, however, would differ from those above only in detail.

The theoretical maximum frequency f_{max} of a taper based acousto-optic device is given by the requirement that the taper waist be at least two-moded. Although the formulae derived above become inaccurate as second mode cutoff is approached, it is possible to estimate the order of magnitude of f_{max} . The acoustic wavelength is given by (9), but the expression (52) for the optical beatlength does not apply. We make use of the fact that the propagation constant of the second mode at cutoff is $\beta_o = kn_{air}$, where $n_{air} = 1$, while that of the fundamental mode is of the order of $\beta_e = kn_o$. Hence the beatlength at cutoff can be approximated by

$$L_B \approx \frac{2\pi}{k(n_o - n_{air})} . \quad (56)$$

Setting $\Lambda = L_B$ for resonance, and eliminating R using the result that, at second-mode cutoff, the V-value is $V_{co} = 2.405$, we obtain

$$f_{max} \approx \frac{(n_o - n_{air})^2}{(n_o^2 - n_{air}^2)^{1/2}} \frac{c_{ext} V_{co}}{2\lambda} . \quad (57)$$

For $\lambda = 633$ nm, (57) gives f_{max} greater than 2 GHz. For $\lambda = 1.3 \mu\text{m}$, f_{max} is around 1 GHz. There is no simple expression for the acoustic power required for maximum conversion at such frequencies, because equations (51) are only accurate far from second mode cutoff. However, (14), (53) and (55) give

$$P \propto \frac{\lambda^{11/3}}{f^{1/3}L^2}, \quad (58)$$

far from cutoff, where R is adjusted to give resonance. Hence the power requirement actually decreases slowly with increasing frequency.

7.3 Bandwidth

The bandwidths Δq are given by (25), with m_q from (26) and Λ / L_B from (53). The result is very simple in this case because, unlike the dual-core fibre, the resonance condition is an analytic function. In fact the dependence of Λ / L_B upon every q can be expressed as $\sim q^{-m_q/2}$, giving for m_q the constants found in Table 2.

m_f	1
m_λ	-2
m_R	3
m_{n_o}	2
m_Λ	-2

Table 2. Values for m_q for a null coupler frequency-shifter.

For the example considered below equation (55), $\Lambda = 0.19$ mm. These results give an RF bandwidth of $\Delta f = 11.6$ kHz, and an optical bandwidth of $\Delta\lambda = 2.45$ nm. A variation in waist diameter of $\Delta(2R) = 0.065$ μm will give a change in resonance wavelength of 10 nm.

7.4 Polarisation

The acousto-optic interaction in a taper waist is expected to be somewhat polarisation-dependent, because there is a large refractive index variation (silica to air) across the waveguide. In the language of optical waveguide theory, the taper waist is not weakly-guiding, unlike the familiar core-cladding waveguide in a single-mode fibre. The taper waist is highly multi-moded, so the polarisation effects are not large and the LP mode approximation is in general valid; nevertheless, they are significant enough to require consideration.

Polarisation dependence arises because each of the modes we have been considering so far, which we have been describing as the fundamental (even) and second (odd) modes, is in fact a pair of modes, each member of the pair being orthogonally polarised with respect to its partner. The electric field distributions of these modes are drawn schematically in Fig. 25. In a circular waveguide, under the LP mode approximation, the two polarisations of the fundamental mode are degenerate. However, the two polarisations of the second mode have slightly different propagation constants. Hence the null coupler is characterised by two slightly different beat lengths, one for each eigen-polarisation. The resonance conditions for the two polarisations will be slightly different, so we can expect the response of the device for unpolarised or arbitrarily-polarised input light to exhibit two resonances.

The two second modes are distinct because one has its electric vector parallel to the nodal line in the field distribution, while the other has the electric vector perpendicular to the nodal line. (In fact there is a further pair of second modes, with field distributions rotated through 90°

compared to "our" modes. These are never excited in the passive null coupler because of symmetry about the plane of the coupler. Any light acousto-optically coupled into them will be lost.)

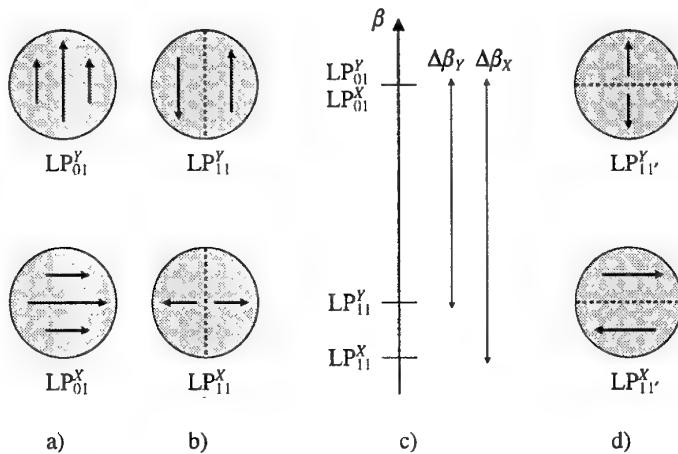


Fig. 25. Schematic diagrams of the amplitude and polarisation of the electric field in (a) the fundamental modes and (b) the second modes of interest. The different propagation constants, and hence beatlengths, for the two states of polarisation are depicted in (c). The other pair of second modes, which are never excited in a low-loss coupler, are drawn in (d).

The splitting of the resonance by polarisation effects can be expressed in terms of the same parameters q as the bandwidth. Taking the results of Ref. 10 for an elliptical waveguide, in the circular limit, the polarisation splitting can be expressed as

$$m_q \frac{\Delta q}{q} = \frac{2j_{11}^2}{(j_{11}^2 - j_{01}^2)} \frac{(n_o^2 - n_{air}^2)}{n_o^2} \frac{1}{V}, \quad (59)$$

where V is the V-value of the cladding-air waveguide at the taper waist.

For our example taper, this gives a polarisation splitting of $\Delta f = 0.08$ MHz in RF frequency and $\Delta\lambda = 16.9$ nm in optical wavelength.

Snyder and Love¹¹ point out that in a perfectly circular waveguide, the LP approximation is inaccurate for the second modes, and the full vector modes must be considered. However, such a treatment gives the same polarisation splitting as (59).

7.5 Longitudinal profile

We have outlined the important equations needed to design a taper or coupler to give acousto-optic resonance for a given acoustic frequency f and optical wavelength λ . This specifies the radius R which the waist must have. It is therefore necessary to be able to know how to make a taper or coupler with a waist of the appropriate radius.

We have stated that the shapes of the taper transitions are unimportant, as long as they provide low-loss and null passive coupling. However, there is a practical reason to be concerned with the diameter profile. The simplest way to make a taper or coupler with a waist of length L is to stretch the fibre(s) in a hot zone of fixed length equal to L . However, for a waist of diameter 6

μm and length 50 mm, and an untreated fibre diameter of 100 μm , this process gives taper transitions which are each 140 mm long. This is impractically long; the fibre(s) would start to sag during taper elongation, and in any case the resultant acousto-optic device would be very cumbersome. We therefore need to consider how to reduce the length of the transitions without shortening the length of the waist.

Complete control of both the waist radius and the entire longitudinal diameter profile of the transitions is possible by changing the length L of the hot-zone during taper elongation⁸. L is given by the travel distance of the oxy-butane burner. In fact a simple two-part variation in L is sufficient to give handlable tapers or couplers with relatively long narrow waists. At first the taper is stretched with a constant hot-zone length of L_o until the taper elongation x is x_1 . Then the hot-zone length is increased linearly with x until $x = x_2$, at which point elongation is stopped. The dependence of L on x can be written as

$$\begin{aligned} L(x) &= L_o & (x < x_1) ; \\ &= L_o + \alpha(x - x_1) & (x > x_1) , \end{aligned} \quad (60)$$

where α is a dimensionless constant between 0 and 1 which determines the rate at which the burner's travel is increased during the second stage. Both L_o and α can be given any convenient values. Smaller values for L_o and larger values for α give shorter transitions, but if the transitions are too short they can cause loss, or a tendency towards non-null coupling. We find that $L_o \approx 10$ mm and $\alpha = \frac{1}{2}$ or $\frac{3}{4}$ are a good compromise.

The critical elongations x_1 and x_2 can be determined using the methods of Ref. 8:

$$\begin{aligned} x_1 &= 2L_o \ln \left\{ \frac{r_o}{R} \left(\frac{L_o}{L} \right)^{1/2\alpha} \right\} ; \\ x_2 &= x_1 + \left(\frac{L - L_o}{\alpha} \right) , \end{aligned} \quad (61)$$

where L and R are the required values for the length and radius of the taper waist, and r_o is the radius of the untreated fibre.

This assumes that it is a single taper which is being made. If a coupler is being made, using two fibres with radii r_1 and r_2 , we must let r_o be the radius of the notional fibre formed by coalescing the two fibres, without elongation, to form a single glass cylinder. By equating areas, we get

$$r_o = \sqrt{(r_1^2 + r_2^2)} , \quad (62)$$

which can be substituted into (61).

8. SINGLE TAPER: EXPERIMENT

8.1 A single taper

To demonstrate the use of the cylindrical taper waist as an acousto-optic medium, without the complication of the need to make a null coupler, simple acousto-optic devices based on the single taper were built. The results from two such devices are described below.

A fibre from the ORC stock (YD392_07, cutoff wavelength 390 nm, diameter 80 μm) was elongated by 84 mm on the coupler rig to form a taper. The throughput loss caused by this treatment was ≤ 0.2 dB. For the first 30 mm of elongation, the burner travel was kept to 10 mm. The travel was then increased by $\frac{3}{4}$ mm for every 1 mm of further elongation. This procedure was designed to produce a taper with a long narrow uniform waist (50 mm long, 6.2 μm diameter), but with relatively short taper transitions (each about 20 mm long) linking the waist to the untreated fibre. The waist-diameter of 6.2 μm was specified to give acousto-optic resonance at 1.5 MHz for 633 nm light.

For efficient acousto-optic coupling, this diameter must be uniform along the entire 50 mm waist, and this is expected from the nature of the travelling-flame elongation method. Nevertheless, uniformity was checked by aiming HeNe light from a fibre end sideways at the taper, and traversing the fibre end along the taper. A diffraction pattern dependent on the local taper diameter was seen behind the taper, and did not vary visibly as the fibre end was traversed.

Despite being very narrow, tapers are surprisingly robust, and easily survive the handling necessary to remove them from the coupler rig to the experimental mount. The taper was suspended, without deliberate tension, between two supports. The fibre was held at each support by being threaded through a section of glass capillary and glued in place. A PZT disc fixed on a rubber platform was placed in contact with a bare but un-narrowed portion of fibre at one end of the taper. Initially, the disc was simply rested against the fibre. Light from a linearly polarised HeNe laser was launched into one end of the fibre, and the light emerging from the other end was detected using a simple reverse-biased photodiode (for broadband dc-coupled operation). The experimental set-up is drawn in Fig. 26, though the horn was omitted here.

The output from an RF signal generator was supplied (via an RF amplifier but without any electrical impedance-matching circuitry) to the PZT disc. The drive frequency was varied until the acousto-optic resonant frequencies for the taper were identified, by observing the resultant drop in throughput light. A PZT disc (resonant frequency 2 MHz, diameter 5 mm) could then be selected which best matched the required resonant frequencies, and this was fixed to the fibre with UV-curing glue for better acoustic coupling into the fibre. Calculation of the expected resonant frequencies of a taper is not accurate enough for deciding which PZT disc to fix to the fibre; the diameter of the uniform taper waist was known only approximately, and the resonant frequency varies (inversely) as the cube of the diameter, through (53). No concentrator horn was used to couple the acoustic energy between disc and fibre.

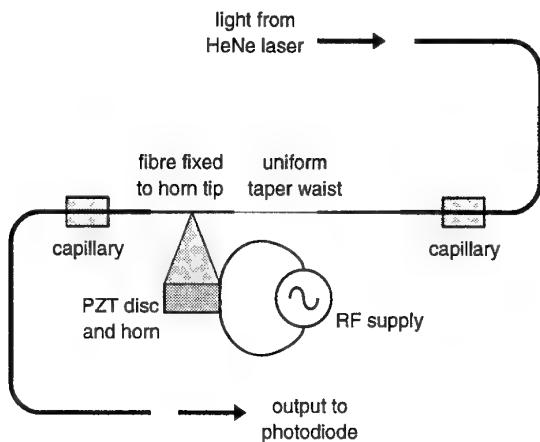


Fig. 26. The assembly of a single taper acousto-optic device.

8.2 Acoustic wave generation

The electrical impedance of the PZT disc used in this experiment were measured using an RF network analyser, for future reference should impedance-matching prove necessary. Fig. 27 is a graph of impedance against frequency; the variation is very rapid.

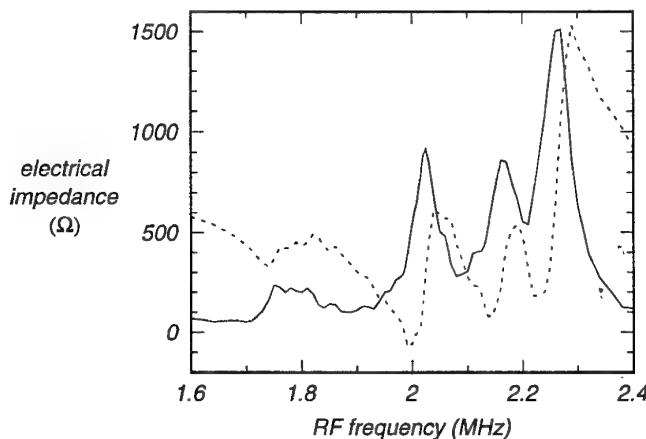


Fig. 27. The frequency dependence of the electrical impedance $Z = R + iX$ of the mounted PZT disc. The solid line is R , and the broken line is $-X$.

The amplitude of the acoustic vibration at the taper waist, the untreated fibre and the disc itself could be measured using the fibre-tip vibrometer described in the earlier report. For a given signal generator output voltage, the vibrometer was used to measure the amplitude of the acoustic wave in the taper for a range of drive frequencies (Fig. 28). As expected, transducer efficiency varied very rapidly with frequency. In fact the variation is so rapid that a reliable calibration curve for amplitude versus frequency would be difficult to produce. When characterising the frequency response of the acousto-optic device, the acoustic amplitude has to be measured afresh using the vibrometer.

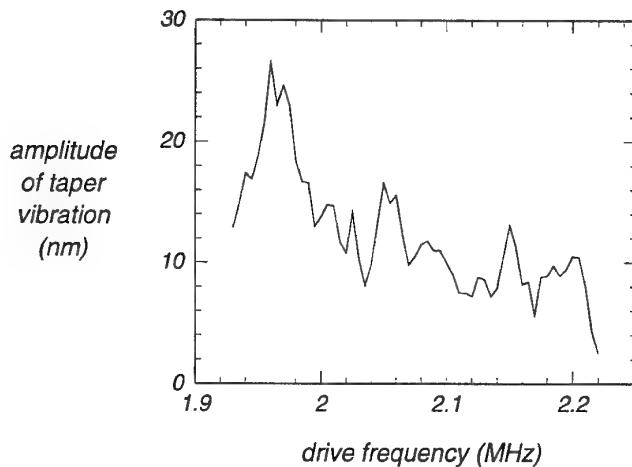


Fig. 28. The frequency dependence of the amplitude of the acoustic wave generated at the taper waist.

For a given frequency, it was found that, as expected, acoustic amplitude is proportional to the RF drive voltage over the range of interest (Fig. 29). This makes it possible to infer values of acoustic amplitude which are too small to measure directly with the vibrometer, by increasing the RF drive voltage until the vibration can be measured, and then scaling the measurement in proportion to the voltage increase. In this way, for a given frequency and drive voltage, it was found that an acoustic wave of amplitude 1.1 nm in the untreated fibre is concentrated by the taper transition to give an amplitude of 26.2 nm in the uniform taper waist. This compares well with the amplitude of 28 nm calculated by assuming the acoustic wave loses no power between the two regions. It also demonstrates that the acoustic transducer does not need to be powerful to give a large amplitude at the interaction region.

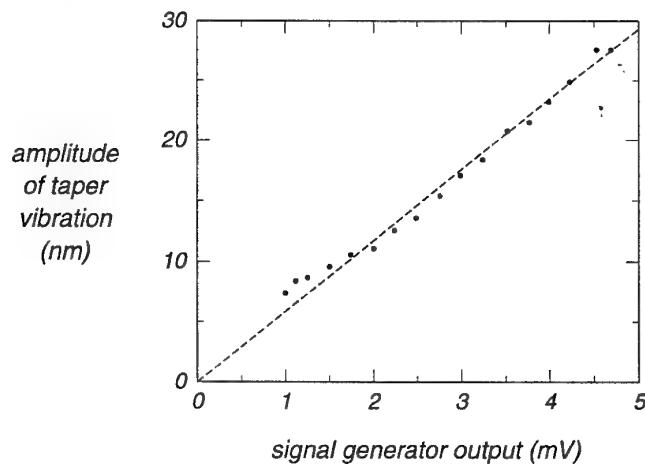


Fig. 29. The amplitude of the acoustic wave generated in the taper waist, as a function of the output of the RF signal generator. This output was amplified before being fed to the PZT disc.

8.3 Acousto-optic resonance

For light with a wavelength of 633 nm, acousto-optic resonance was expected for frequencies around 1.5 MHz. The signal generator frequency was varied about this value until dips in the

optical throughput, corresponding to acousto-optic resonances, were seen. Numerous resonances were observed within the frequency range 2.086 to 2.155 MHz, with a weaker set in the range 1.931 to 2.013 MHz. An optical throughput reduction of 95 % was possible for an applied RF amplitude of 2.4 V pk-pk at the frequency of 2.124 MHz. The required amplitude of the vibration in the uniform taper waist was 9.5 nm, corresponding to an acoustic power of 260 nW. A further increase in drive voltage (and hence acoustic amplitude) brought about a rise in throughput, indicating that the maximum conversion of light from the fundamental to the second mode was 95 %. This is shown, for a different resonance, in Fig. 30.

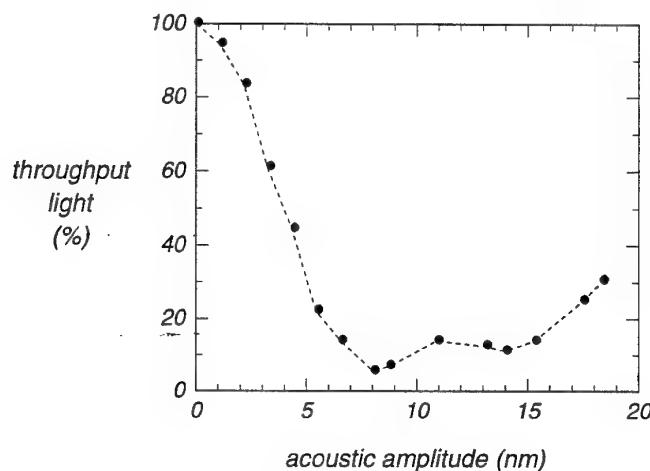


Fig. 30. The optical throughput versus the amplitude of the acoustic wave at the taper waist, for $\lambda = 633 \text{ nm}$ and $f = 2.1436 \text{ MHz}$.

The expected value of acoustic amplitude for maximum conversion is 3.7 nm (corresponding to an acoustic power of about 40 nW), and the expected maximum conversion is of course 100 %. The higher than expected acoustic amplitude (and hence much greater acoustic power, since power varies as the square of the amplitude) probably has the same cause as the distribution of minor resonances within the two RF frequency bands, namely some non-uniformity in the diameter of the taper waist, giving a non-uniformity in the local resonant frequency. A variation in diameter of just 1 % is sufficient to account for the frequency variations above. The displacement of the resonant frequencies upwards from the expected value of 1.5 MHz is assumed to be due to an inaccuracy in the determination of the mean waist diameter during taper elongation.

The less-than-ideal maximum conversion (together with the existence of two ranges of resonant frequencies) is probably caused by polarisation effects. If the drive frequency is on-resonance for one eigen-polarisation but off-resonance for the other, only light in the first polarisation will be converted to the second mode. Although it is likely that there will be more of one eigen-polarisation present than the other (since the output of the HeNe laser is polarised), there will nevertheless be some of both. Hence a measured maximum conversion less than 100 % is expected in the absence of polarisation control. At some of the resonant frequencies, it was found that the maximum conversion depended strongly on the angle of the plane of polarisation of the input light wave. This is as would be expected if each resonance coupled light of one polarisation only. Also, the observed separation of the two frequency ranges is close to the polarisation splitting of 0.11 MHz predicted by theory. However, our measurements of polarisation dependence are merely suggestive at this stage, since a plane polarised wave launched into the fibre will suffer unpredictable changes in polarisation state, due to birefringences in the fibre pigtail.

8.4 Acoustic reflection

When an RF signal at a resonance is supplied to the PZT disc, the throughput light level falls as described above. On an oscilloscope trace of the photodiode output, there is also a small time-varying modulation in the throughput, with a frequency twice that of the RF drive. Acoustic reflection at the opposite end of the taper to the PZT disc generates an image sideband signal. This has a beat frequency of $2f$ with the principal frequency-shifted wave.

For a given acoustic reflection coefficient, the maximum modulation occurs when there is 50 % conversion between the fundamental mode and the second mode in the taper. As a fraction of the throughput light level with no RF drive, the pk-pk modulation varies between 2.5 % and 6 % depending on the frequency. This corresponds to an acoustic power reflection coefficient less than 0.05 % from (31).

Measurements were made of the acoustic amplitude at various points along the taper, over a length of a few acoustic wavelengths. Any standing wave component was too small to be discerned above the measurement uncertainty of the vibrometer.

8.5 Stability and drift

The resonant frequencies, amplitudes and efficiencies of the device appeared to drift over a time-scale of minutes. This precluded any quantitative broad-band measurements of device performance as a function of frequency, because each measurement (which would include vibrometry, an exploration of the variation with the polarisation state of the input light, and a very small frequency step between measurements) would take too long. This drift is likely to be due to the environmental sensitivity of the experimental holding arrangement.

More deliberate disturbances, such as a change in the tension in the taper, also brought about a change in the device's resonant frequencies, and other characteristics. It is likely that such disturbances slightly change the nature of the waveguide at the uniform taper waist, and also perturb the acoustic dispersion relation, thus changing the conditions of acousto-optic resonance. As well as being a nuisance, the strong sensitivity of the taper to its environment could hopefully be used as a tuning mechanism.

8.6 Parasitic resonances

One advantage of acousto-optic devices based on tapered fibre is the low RF drive power required. This makes possible overcoupled operation over short interaction lengths, without taxing the efficiency of the transducer. However, an exploration of overcoupled operation in the present device, demonstrating that it could be shortened, was not possible. For acoustic amplitudes greater than that required to give maximum conversion (typically between 15 and 30 nm), a set of parasitic vibrations become excited along the fibre. These could be picked up by the vibrometer, and also disturbed the optical throughput signal of the device. On the RF spectrum analyser at the vibrometer output, this vibration exhibits two characteristics; a comb of peaks separated by ~160 kHz and a modulation at half the drive frequency (Fig. 31). The parasitic vibration is strong and commences suddenly as transducer amplitude is increased above a threshold. It is evident on the un-narrowed fibre, though to a lesser extent than on the uniform taper waist (Fig. 31e). This suggests that it is caused by unwanted modes of vibration

of the fibre where it is attached to the PZT disc. The possibility that the vibrating taper was hitting the vibrometer fibre's tip was discounted, because the resonances were visible in the throughput signal without the vibrometer in place, and could be measured by the vibrometer even when it was clearly separated from the taper.

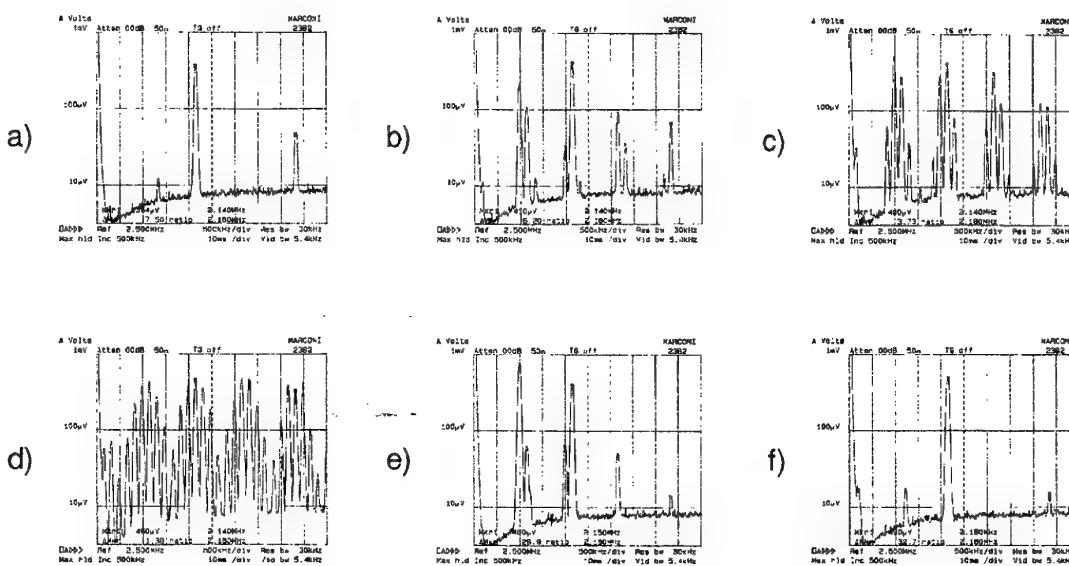


Fig. 31. Frequency spectrum of vibrometer output for an RF drive at 2.1525: (a)-(d) at the uniform taper waist, with RF amplitudes of 3.5, 4.4, 7.3 and 17.5 V respectively; (e) at the untreated fibre (17.5 V); and (f) on the PZT disc (17.5 V). The peak at 1.33 MHz is an artefact caused by the vibrometer laser.

8.7 Another single taper

The second taper was elongated with a burner travel of 10 mm until the taper extension was 20 mm. Then the burner travel was increased by 1 mm for every 2 mm of further extension. Taper elongation was stopped when an extension of 100 mm was reached. This gave a taper waist 50 mm long, with more gradual transitions each 30 mm long.

8.8 Acousto-optic resonance

A PZT disc with a resonant frequency of 1.5 MHz was attached to the untreated fibre, this time via an aluminium concentrator horn (Fig. 26). For $\lambda = 633$ nm, the strongest acousto-optic resonance was centred at $f = 1.936$ MHz.

As mentioned before, the amplitude of vibration of the PZT disc varies very strongly with frequency. However, some indication of the device's frequency response could be obtained by stepping through a range of frequencies. At each frequency, the RF drive voltage was adjusted until the vibrometer indicated a certain acoustic amplitude (~12 nm). A half-wave plate at the output of the HeNe laser permitted some limited control over polarisation. Graphs of device throughput against drive frequency are drawn in Fig. 32 for each of two wave plate angles; polarisation-dependent behaviour, with a splitting of 0.08 MHz, is indicated. This compares with the theoretical polarisation splitting of 0.10 MHz.

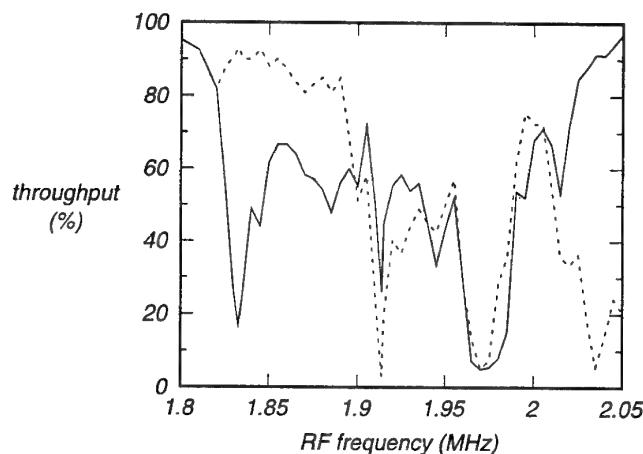


Fig. 32. The frequency dependence of the optical throughput, for input light in each of two orthogonal polarisation states.

8.9 Wavelength dependence

The dependence of device output with optical wavelength was measured by launching white light into the device, and measuring the output spectrum using an optical spectrum analyser. Fig. 33a shows the wavelength dependence of the throughput for a drive frequency of 2.05 MHz, with the drive voltage adjusted to give maximum conversion. The FWHM bandwidth of ~ 50 nm is much greater than the theoretical bandwidth of 2.60 nm.

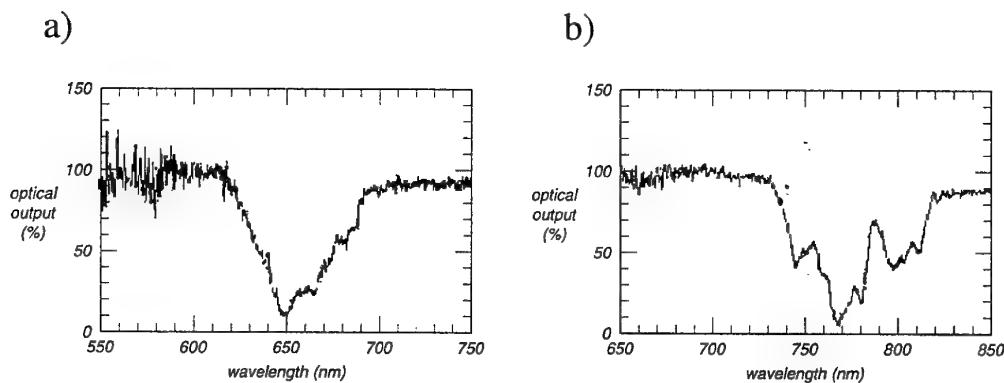


Fig. 33. Output spectra for $f = 2.05$ MHz, with the taper (a) taut and (b) slack.

The characteristics of the device are dependent on the tension in the taper. When the taper was made slack, the resonance wavelength changed by about 100 nm for a given acoustic frequency. The spectrum also broke up to give a three-peaked structure, as shown in Fig. 33b. Without a quantitative measurement of the tension or the strain of the fibre, it was not possible to compare this result with the theory embodied in (34).

The explanation for a three-fold resonance probably rests with the polarisation properties of the complete vector set of four second modes¹¹, though the details are not clear as yet. There would be three different beatlengths for mode coupling between fundamental and second modes, one

of which will involve coupling to a degenerate pair. This interpretation is supported by the splitting of 20-25 nm between the peaks, which is close to the expected polarisation splitting of 18 nm. The FWHM bandwidth of ~50 nm is then composed of distinct peaks with FWHM of ~10 nm, separated from each other by the polarisation splitting.

The centre wavelength of the resonance changes when the drive frequency is changed. The variation of resonance wavelength with frequency for the taper is drawn in Fig. 34, both for the slack and for the original taut taper.

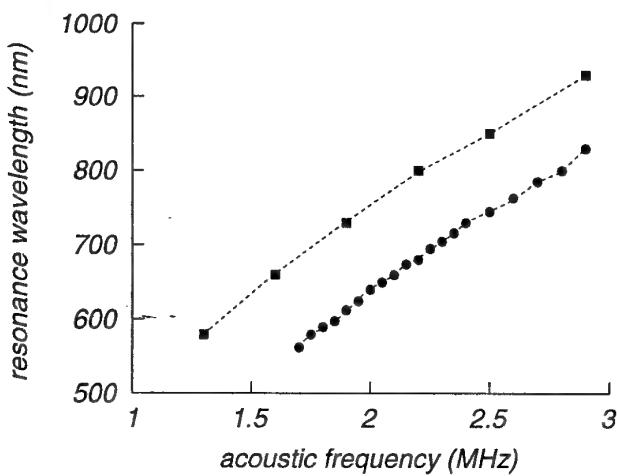


Fig. 34. The frequency dependence of the resonance centre wavelength, when the taper is taut (●) and slack (■).

8.10 Power requirements

The electrical drive power required to give maximum conversion is an important parameter to measure. There was no RF power meter at the ORC capable of measuring the small drive power of the device, so the power was calculated from measurements of the RF current and voltage supplied to the PZT disc, together with the phase angle between them.

The device was adjusted to give maximum conversion at a strong resonance. The acoustic amplitude required was measured using the vibrometer to be 21.5 nm. The PZT disc was connected in series with a known resistor $R = 10 \Omega$ using the circuit of Fig. 35. The two voltage outputs V_1 and V_2 were connected to x and y inputs of an oscilloscope with the time-base switched off. The amplitude and phase of the current and voltage supplied to the series combination can be determined by examination of the elliptical trace on the oscilloscope screen, and hence the power consumed by the combination can be calculated. The power dissipated by the resistor alone is readily calculated and subtracted to leave the drive power of the PZT disc. This was calculated to be 0.50 ± 0.05 mW, much smaller than the smallest published value for the drive power of any other acousto-optic fibre device.

This value can be compared to the acoustic power excited on the taper, as calculated from the amplitude of 730 nW measured with the vibrometer. The theoretical acoustic power required is 42.4 nW. This indicates that the small measured RF drive power is in fact very much larger than it could be, and could be reduced if the transducer efficiency were optimised by electrical impedance-matching and/or more sophisticated coupling of acoustic energy to the fibre.

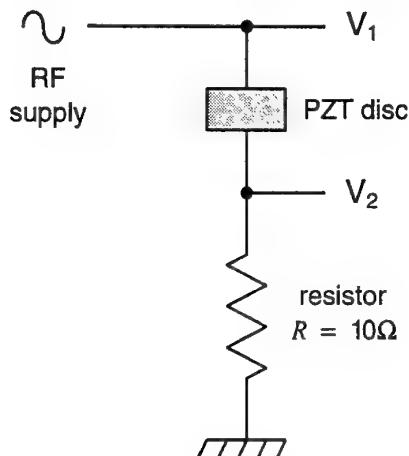


Fig. 35. The circuit for measuring RF drive power.

8.11 Parasitic resonances

The second taper also exhibited parasitic resonances, although their onset was postponed by the more efficient acoustic wave generation afforded by the presence of the concentrator horn. These resonances still have no satisfactory explanation, but it was possible to gather a further piece of evidence by measuring the dependence of the parasitic frequency on the drive frequency, which is plotted in Fig. 36.

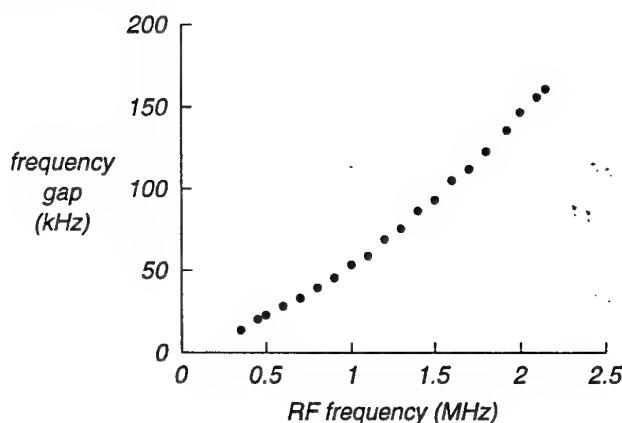


Fig. 36. The characteristic frequency of the parasitic resonance versus the RF drive frequency.

9. NULL COUPLER: EXPERIMENT

9.1 Null coupler

The first step in making the frequency shifter was to make a null fused tapered coupler with a circularly-fused waist cross-section. There is no known report of this ever having been done before, presumably because a null coupler is, on the face of it, a singularly useless component. The requirement that the waist be circularly fused makes it more difficult to ensure zero splitting. A range of dissimilar fibres was tried, and eventually a pair of fibres was identified which readily yielded circularly-fused null couplers. These fibres were ND281_02 and ND313_03, with diameters $80 \mu\text{m}$ and $60 \mu\text{m}$, NAs of 0.21 and 0.26, and cutoff wavelengths of 650 nm and 500 nm, respectively (note that the cut-off wavelengths quoted for these fibres in the 1993 Annual Report are incorrect). Unfortunately, ND281_02 is not single-mode at $\lambda = 633 \text{ nm}$, but by always launching into the other fibre there was no danger of this causing problems. If it were ever necessary to launch into ND281_02, second-mode strippers could always be placed before and after the coupler.

Values of maximum intrinsic coupling M less than 1:6000 were possible with these fibres. Indeed, the small amount of light coupled into the second fibre appeared always to be in the second mode of the fibre. By stripping the second mode, splitting ratios even closer to the ideal of zero should be possible. One of these couplers, with HeNe light at 633 nm launched into the fibre ND313_03, was cleaved at its waist, and the light emerging from the cleaved waist could be examined in the far field by projection onto a screen. The far field was (to the eye) a perfect two-lobed second-mode pattern, as shown in Fig. 37. Since the launch fibre was the fibre with the lesser cutoff wavelength, it is expected that light in its fundamental mode should evolve into the second mode of the coupler waist. The observed far field pattern indicates that the coupler was indeed a circularly-fused null coupler. Circular fusion was finally confirmed by examining the cleaved coupler end under a SEM. Fig. 38 is a photomicrograph showing a circular waist, with a diameter of $5.5 \mu\text{m}$.

The null coupler for a frequency-shifter experiment was made by elongating the parallel pair of fibres with a burner travel of 10 mm until the taper extension was 36 mm. Then the burner travel was increased by 1 mm for every 2 mm of further extension. Coupler elongation was stopped when an extension of 66 mm was reached. This gave a taper waist 25 mm long, with an expected diameter of $6.5 \mu\text{m}$ for acousto-optic resonance at $f = 1.25 \text{ MHz}$. The coupler's excess loss was $\sim 0.1 \text{ dB}$. The maximum splitting ratio was relatively large at 1:400, though again the coupled light was mainly in the second mode of the fibre, and therefore could have been stripped away if necessary.

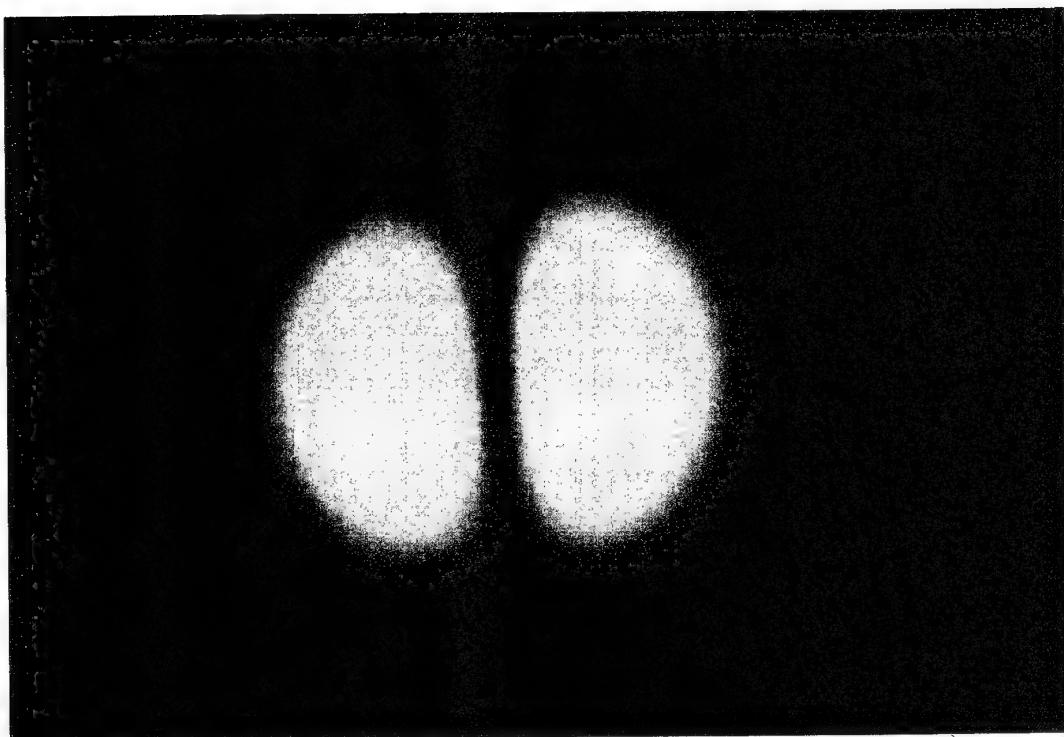


Fig. 37. A photograph of the far field pattern emerging from the cleaved waist of a null coupler.

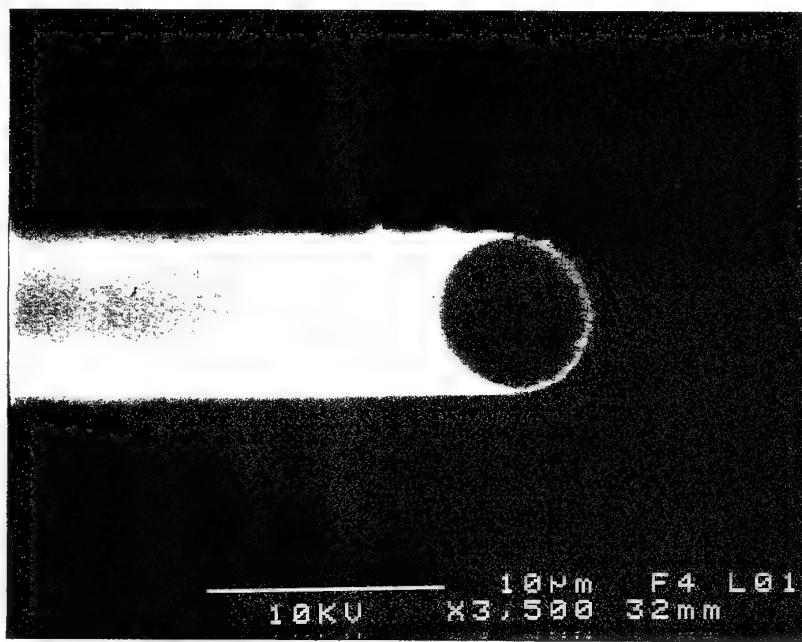


Fig. 38. SEM of the cleaved waist of a null coupler.

Unlike the single taper, the coupler must be mounted so that the acoustic wave vibrates in a particular plane, namely the plane of reflection symmetry of the coupler. Small paper tags were fixed to the parallel untreated fibre pairs at either end of the coupler to assist in alignment. The coupler was mounted on the same assembly as the tapers had been. However, the PZT disc was replaced with a large 12 mm square PZT plate with a resonant frequency of 1.25 MHz, mounted firmly on a solid aluminium base instead of a rubber base. An aluminium cone with a base diameter of 12 mm was fixed onto the plate. The experiment was set up so that the acoustic wave and optical waves propagated in opposite directions along the fibre.

9.2 Acousto-optic resonance

Resonances were found in two distinct frequency ranges, about 1.79 MHz and 1.86 MHz. For arbitrary input polarisation, the graph of optical output in the throughput fibre against signal generator output is similar to that for the single taper experiments, except this time the converted light is guided in the second fibre instead of being lost. Although this fibre was known to be two-moded, the coupled light was carried in the fundamental mode.

A fibre polarisation controller and a bulk half-wave plate were set up on the input fibre. Although the controller was unsatisfactory and gave only partial control of polarisation, it could be adjusted to give a maximum conversion greater than 99 % at some frequencies. Fig. 39a is a graph of the output optical powers against signal generator output at 1.8510 MHz, with the polarisation controller adjusted to give a maximum conversion of about 98 %. Outputs are also plotted for the orthogonal polarisation, obtained by rotating the half-wave plate through 45°. This resonance is clearly very polarisation dependent. Without re-adjusting the polarisation controller, the drive frequency was changed to 1.7952 MHz, in the other range of resonances, and the measurements were repeated. The results are drawn in Fig. 39b. Although not so clear cut, this demonstrates that the presence of two frequency ranges corresponds to polarisation splitting of the acousto-optic resonance. As confirmation, the calculated polarisation splitting of 0.08 MHz is close to the measured splitting.

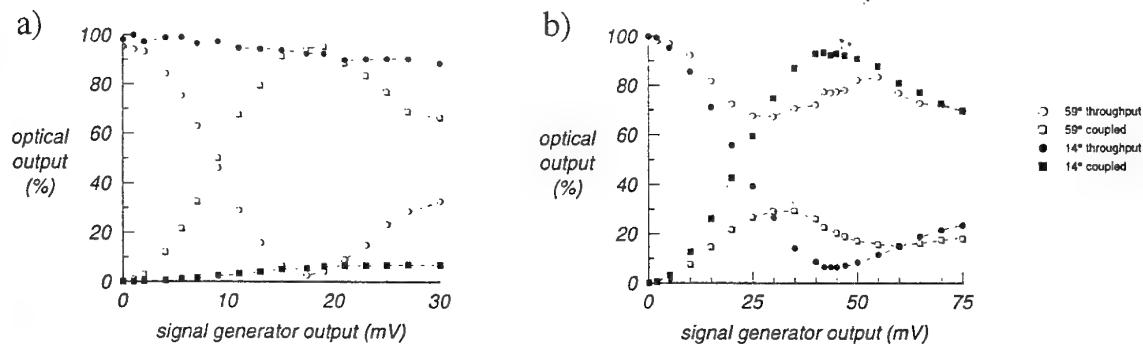


Fig. 39. The optical output powers in each fibre, for each of two orthogonal input polarisations, for drive frequencies of (a) 1.8510 MHz, and (b) 1.7952 MHz.

9.3 Wavelength dependence

The optical spectra of the outputs of the device were measured in the same way as with the single taper. A typical pair of spectra, for a drive frequency of 1.860 MHz, is drawn in Fig. 40 for both output fibres. The throughput output spectrum shows three dips, while the coupled

spectrum shows three peaks at corresponding wavelengths. The spacing of 15 nm between the peaks is close to the expected polarisation splitting of 18 nm, leaving an intrinsic FWHM for each resonance of ~10 nm. Again, this is greater than expected, presumably because of slight non-uniformities in the diameter of the taper waist.

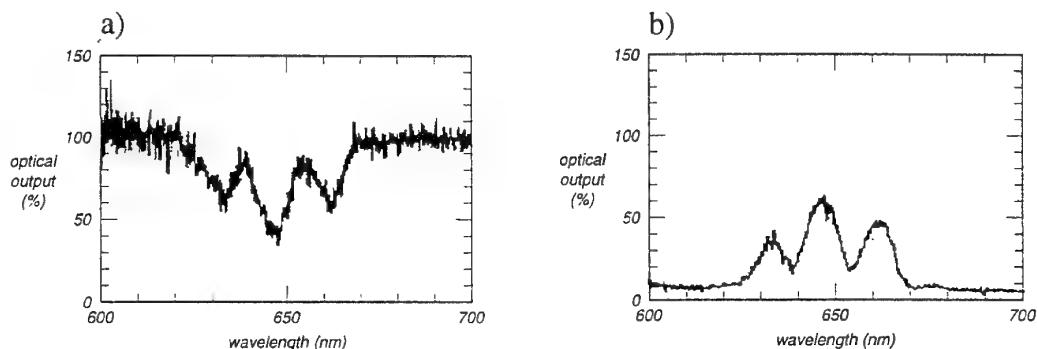


Fig. 40. Output spectra of (a) the throughput light and (b) the coupled light, for $f = 1.860 \text{ MHz}$.

As in the single taper experiment, a three-fold resonance was observed. However, under some circumstances the central resonance can split into two, giving a total of four clear resonances. Fig. 41 is the throughput spectrum when the device was driven at 2.5 MHz. The presence of three or four strong resonances in the output spectra is not fully understood. Clearly it is related to the vector polarisation properties of the second mode in the coupler waist. For a circular waveguide, there are four second-modes, two of which are degenerate¹¹. The degeneracy between these two can be lifted by a slight ellipticity, or maybe by the effect of the residual fibre cores. This gives three or four possible values for the beatlength between fundamental and second modes, and hence that number of resonances.

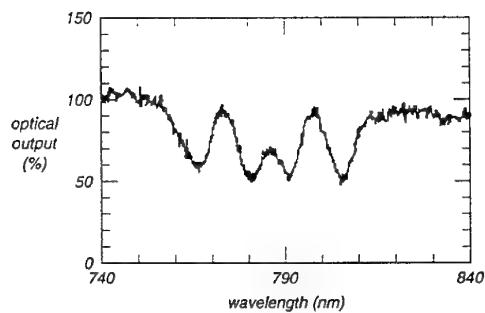


Fig. 41. Output spectrum of the throughput light for $f = 2.5 \text{ MHz}$.

This type of explanation is unsatisfactory for two reasons. Firstly, we know that our input wave is exciting the second mode of the coupler waist, not the fundamental. Hence three or four resonances would only be seen if three or four second-modes were excited. However, we can only excite a maximum of two modes without there being some mode-coupling, because there are only two input modes to launch into (namely the two orthogonally polarised fundamental modes of the input fibre). We know there is no significant mode-coupling, or else the passive null coupler would exhibit a wavelength-dependent response. Secondly, we know that two of the second modes of the coupler waist do not correspond to any light guided in either fibre, but instead correspond to cladding modes which are lost. However, all three or four observed resonances are in fact retained in the output fibres.

9.4 Drive power

This was measured in the same way as for the single-taper experiment. For maximum conversion at 1.857 MHz, the drive power was ~ 1 mW. Greater values for the drive power were measured at other frequencies. The corresponding acoustic amplitude was measured to be $u_o = 27$ nm using the vibrometer. The calculated actual acoustic power of $2.0 \mu\text{W}$ can be compared with the electrical drive power, and with the theoretical required acoustic power of 170 nW. The comparisons indicate that the efficiency of acoustic wave generation is very low, and that the coupler waist is not perfectly uniform.

9.5 Frequency shift

It is inconceivable that the light emerging from the coupled port of the device could be unshifted in frequency. However, it was necessary to confirm this, and also measure how pure the frequency-shifted wave was. A Mach-Zehnder interferometer was set up as shown in Fig. 42, using bulk optical components. The frequency shifter was inserted into one arm of the interferometer, with the output taken in turn from both fibres. A conventional Bragg cell was inserted into the other arm. The Bragg cell was driven at 80 MHz, up-shifting the frequency of the light wave in that arm of the interferometer. The output of the detector was fed into an RF spectrum analyser, which displayed the frequency spectrum of the beat signal between the 80 MHz upshifted light and each output of the device.

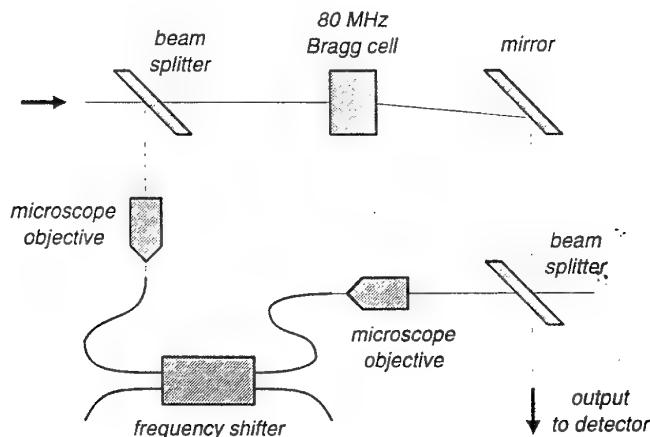


Fig. 42. Measurement of the frequency shift.

For a strong resonance at 1.857 MHz, with the drive amplitude adjusted to give a maximum conversion of over 96 %, the frequency spectrum of Fig. 43 was obtained for the coupled output. The main frequency component present was at about 81.80 MHz. This corresponds to a wave which is down-shifted in frequency (so that the beat frequency with +80 MHz is greater than 80 MHz). This is as expected assuming coupling is from second mode to fundamental mode by a counter-propagating acoustic wave. Also visible are signals at the carrier frequency (79.96 MHz) and the image sideband (78.12 MHz). These are both about 30 dB down from the principal signal.

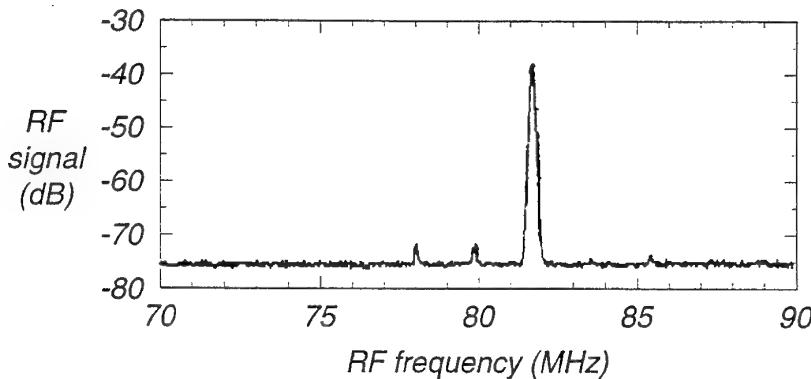


Fig. 43. The RF spectrum of the output signal of the detector in the heterodyne experiment, where light from the coupled output of the frequency shifter is interfered with light which is frequency up-shifted by ~80 MHz in a Bragg cell.

The proportions of shifted, carrier and image sideband power remain roughly the same as the RF drive voltage is reduced from the value for maximum conversion. Fig. 44 is the spectrum when the drive voltage was adjusted to give 50 % conversion. Also drawn is the corresponding spectrum for the output in the throughput fibre, with a mainly unshifted signal as expected.

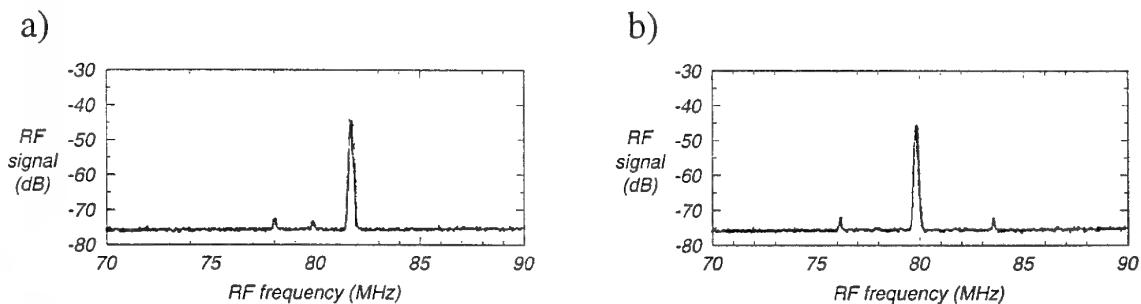


Fig. 44. RF spectra at 50 % conversion for the (a) coupled and (b) throughput outputs.

The light at the carrier frequency contaminating the frequency-shifted output is due to the small amount of light coupled by the passive coupler; that is, its deviation from true null behaviour. We anticipate that the measured carrier suppression can be readily improved because (a) better null couplers have been made with splitting ratios over ten times smaller, and (b) most of the passively coupled light was in the second mode of the fibre, and could therefore be stripped easily.

The light in the image sideband could be generated by deviation of the coupler from ideal null behaviour, but on its own the splitting ratio of 1:400 would give an image sideband suppression better than 50 dB, so this is not a concern. The most likely source of the image sideband component we observe is the presence of a reflected acoustic wave in the coupler waist.

9.6 Repeatability

Since the device described above was made, further null coupler frequency shifters have been made from various fibres and operating at various optical wavelengths and acoustic frequencies. All exhibited similar behaviour to the prototype, indicating that the frequency shifter is reproducible.

10 NULL COUPLER USING IDENTICAL STANDARD FIBRES

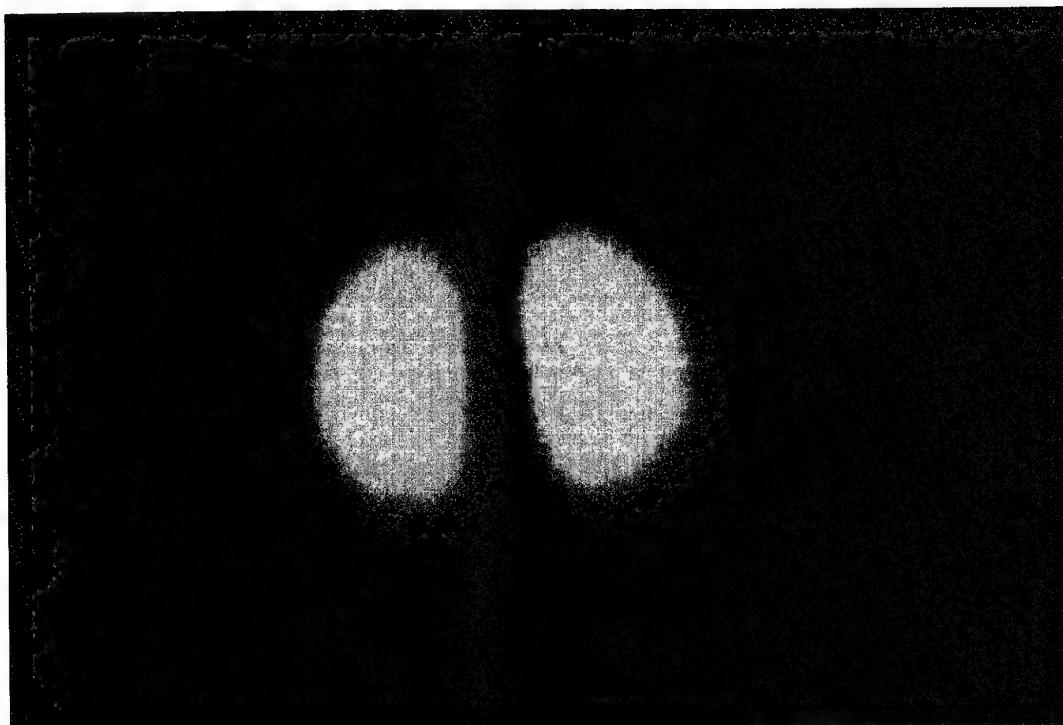
The most pressing aim of this project is to realise an acousto-optic frequency shifter operating at a standard wavelength using standard single mode telecommunications fibre. The prototype frequency shifter was made using dissimilar fibres. We first address the problem of making a null coupler with identical ports, then we make frequency shifters using such a null coupler made from a standard single-wavelength telecommunications fibre.

10.1 Null coupler using common fibre

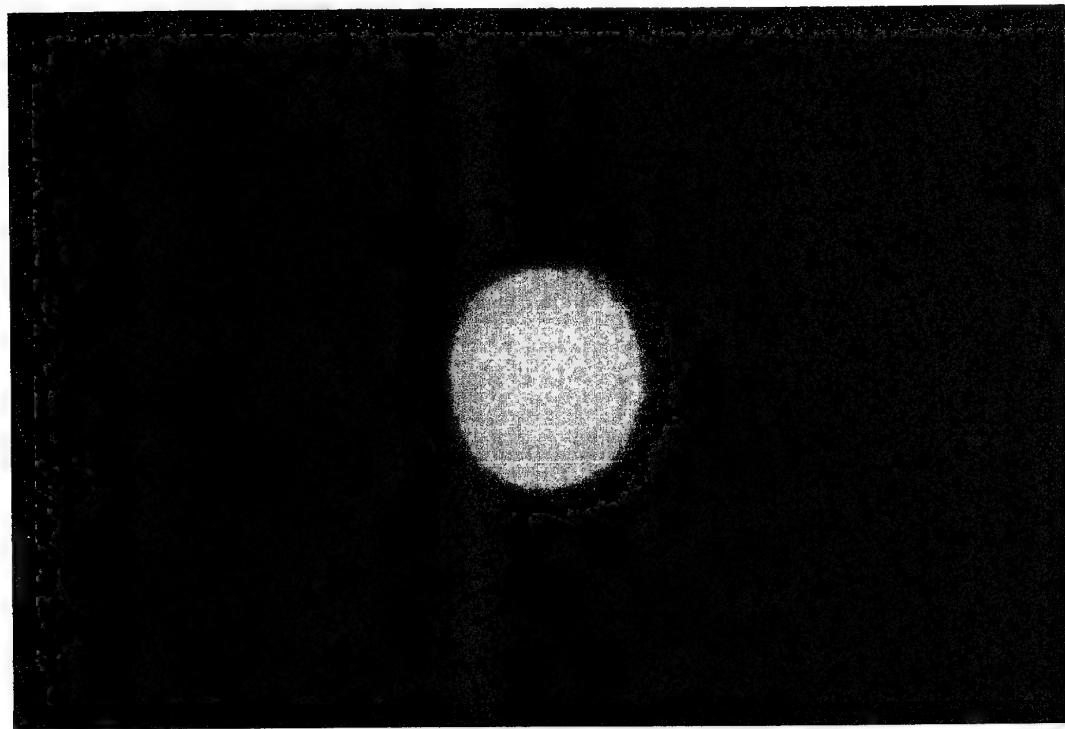
The prototype was made from two dissimilar fibres, one of which was two-moded at the operating wavelength of 633 nm. To make a device with all single-mode fibres, all that is necessary is to select two dissimilar single-mode fibres which will yield a null coupler. However, to give identical ports, it is first necessary to make a pair of identical fibres dissimilar over a short length. The coupler can then be made within the short length, leaving identical ports. The means for making fibres dissimilar prior to coupler fabrication are well known, and are used in making "wavelength-flattened" couplers. The favoured method is pretapering, in which one of the fibres is stretched slightly on its own, before being stretched together with the other fibre to form a coupler¹³. This pre-treatment method has the advantage of requiring the same equipment as is needed for making the coupler itself. Other published methods include pre-etching, in which some cladding material is chemically removed from one fibre using HF, and pre-polishing, in which a segment of the cladding is mechanically abraded.

To implement this approach, we started with one type of fibre, and made a series of fused taper couplers, in each of which one of the constituent fibres had been pretapered to a different extent. We could thus determine the pretaper ratio necessary to give a satisfactory null coupler. The fibre used was ND 132_09. This fibre has an outer diameter 75 μm and a cut-off wavelength of less than 450 nm. It was therefore single-mode at 633 nm.

Because of the previous success using dissimilar fibres, with maximum splitting ratios <1:400 commonly being achieved, it was decided to pretaper one length of fibre from 75 μm to 56 μm . The diameter ratio of the two fibres in the region to be fused would then be the same as before, i.e. 3:4. To achieve this, the fibre was elongated through 11.5 mm with a burner travel of 20 mm. The pretaper was accompanied by very little loss, <0.08 dB. The null coupler was then made by elongating the parallel pair of fibres with a burner travel of 10 mm until the taper extension was 31 mm. The burner travel was increased by 1 mm for every 2 mm of further extension. Coupler elongation was stopped when an extension of 61 mm was reached. This gave a taper waist 25 mm long, with an expected diameter of 8 μm . The coupler excess loss was ~0.3 dB, and its maximum splitting ratio was 1:600. To illustrate the excitation of the modes in the coupler waist, the waist was cleaved and the light emerging examined in the far field by projection on to a screen. The far field patterns when laser light is launched into the pretapered fibre, and then the un-pretapered fibre are shown in Fig. 45. The two lobed pattern illustrated by (a) confirms that light in the pretapered fibre excites the second mode in the waist, whereas light in the un-pretapered excites the fundamental mode as shown in (b).



(a)



(b)

Fig. 45. A photograph of the far field pattern emerging from the cleaved waist when light is launched into (a) the pretapered and (b) the un-pretapered fibre.

10.2 Frequency shifter at 1550 nm using standard single-mode telecom fibre

Having shown that null couplers can be made using identical fibres, the procedure of the previous section was repeated with standard single-mode telecommunications fibre. A number of devices was assembled. The two frequency shifting devices discussed below were operated at frequencies of 500 kHz and 12 MHz. Both devices were fabricated using 125 μm fibre with a cutoff wavelength of < 1300 nm, manufactured by Optical Fibres Deeside, and a pigtailed DFB laser at $\lambda = 1550$ nm was used as the source.

10.3 Frequency shifter at 500 kHz

A fibre frequency shifter was made for 500 kHz. The null coupler was made as described in section 10.1. One fibre was pretapered from 125 μm to 94 μm , again satisfying the diameter ratio of 3:4 for the pretapered and un-pretapered fibres. The loss incurred by pretapering the fibre was too small to be measured (<0.01 dB). For a frequency shift of 500 kHz at $\lambda = 1550$ nm a waist diameter of 16 μm was required for the null coupler. The excess loss of the coupler was 0.03 dB and the maximum splitting ratio was 1:2000.

10.3.1 Acousto-optic resonance

The null coupler was mounted on the same assembly as for the prototype of section 9, with the acoustic wave vibrating in the plane of reflection symmetry of the coupler. The PZT was firmly cemented to an aluminium base and had a resonant frequency of 1.25 MHz. An aluminium cone with a base of diameter 12 mm was fixed to the PZT. The experiment was set up so that the acoustic and optical waves propagated in opposite directions.

A fibre polarisation controller was set up for the input fibre. Maximum conversion of ~98% was obtained for an acoustic frequency of 571 kHz when the controller was suitably adjusted. Fig. 46 is a graph of the frequency shifter outputs as a function of pk-pk RF drive voltage across the PZT transducer. The light was launched into the un-pretapered fibre.

10.3.2 Wavelength dependence

The optical spectra of the outputs were measured in the same way as for the prototype. With a drive frequency of 571 kHz, a white light source was launched into either arm in turn. The throughput and coupled spectra were recorded as shown in Figs. 47 and 48. As for the prototype, a threefold resonance is observed with a spacing of 20 nm between the peaks. Again this is close to the expected polarisation splitting of 18 nm, leaving an intrinsic FWHM for each resonance of ~15 nm.

10.3.3 Drive power

This was measured in the same way as for the prototype. For maximum conversion at 571 kHz, the drive power was ~2 mW. The corresponding acoustic amplitude was measured as $u_0 \approx 85$ nm using the vibrometer. The actual acoustic power calculated from this amplitude is 8.47 μW , while the theoretical required acoustic power is 6.6 μW . The comparison between the actual and theoretical acoustic powers indicates that the coupler waist is uniform, though clearly the efficiencies of acoustic wave generation from the RF drive is still very low.

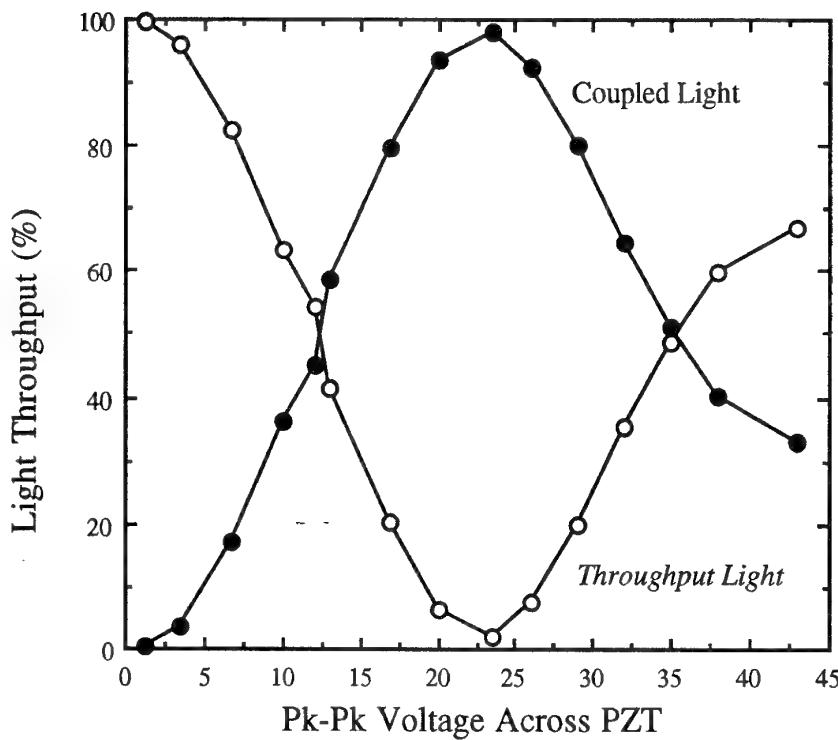
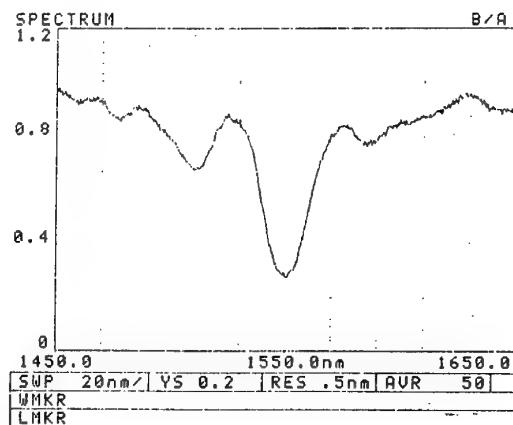


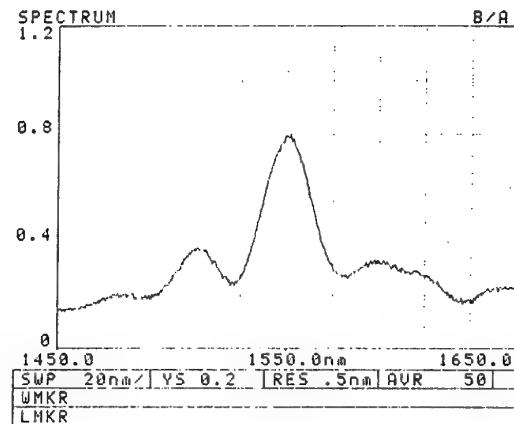
Fig. 46. Performance of the frequency shifter as a function of pk-pk drive voltage across the PZT transducer.

10.3.4 Frequency shift

In order to measure the frequency shift, an all-fibre Mach-Zehnder interferometer was set up as shown in Fig. 49. This configuration was preferred to the version of Fig. 42 because the laser source was pigtailed. The frequency shifter was spliced into one arm of the interferometer. A conventional Bragg cell (80 MHz) was inserted into the opposing arm, so the laser output from this arm first had to be collimated before being upshifted. The output from the Bragg cell was then relaunched into the second input port of the second coupler. The heterodyned signal was detected at one of the two outputs from the second coupler. Unused coupler ports were immersed in index matching fluid to minimise reflections. The beat signal between the active frequency shifter and the Bragg cell is shown in Fig. 50. The main frequency component present is at 80.56 MHz. This corresponds to a wave downshifted in frequency. This is expected assuming coupling is from the second mode to fundamental mode by a counter-propagating acoustic wave. Also visible are signals at the carrier frequency (79.99 MHz) and the image sideband (79.42 MHz). The carrier signal is typically down by 30 dB from the principal signal, but the image side band is surprisingly large, only 20 dB down.

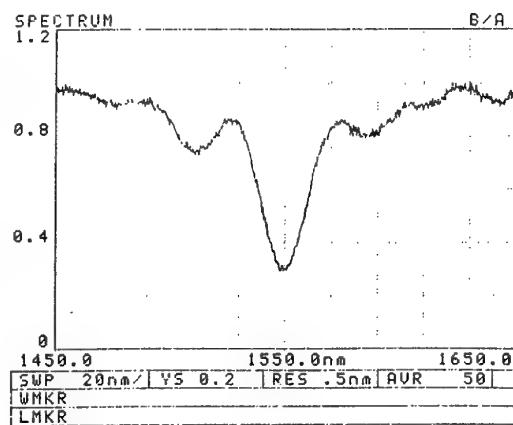


(a)

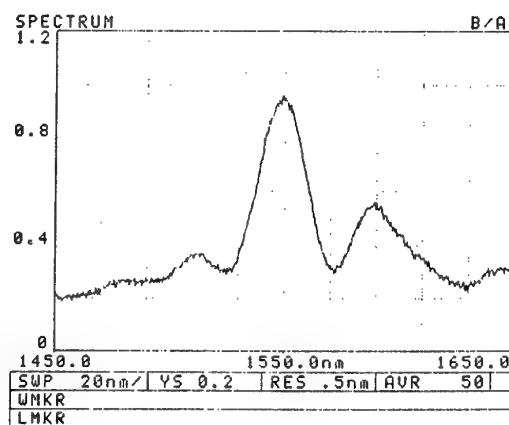


(b)

Fig. 47. Output spectra for light launched into the pretapered arm: (a) the throughput and (b) the coupled light, for $f = 571$ kHz. Wavelength (nm) is given on the horizontal scale.



(a)



(b)

Fig. 48. Output spectra for light launched into the un-pretapered arm: (a) the throughput and (b) the coupled light.

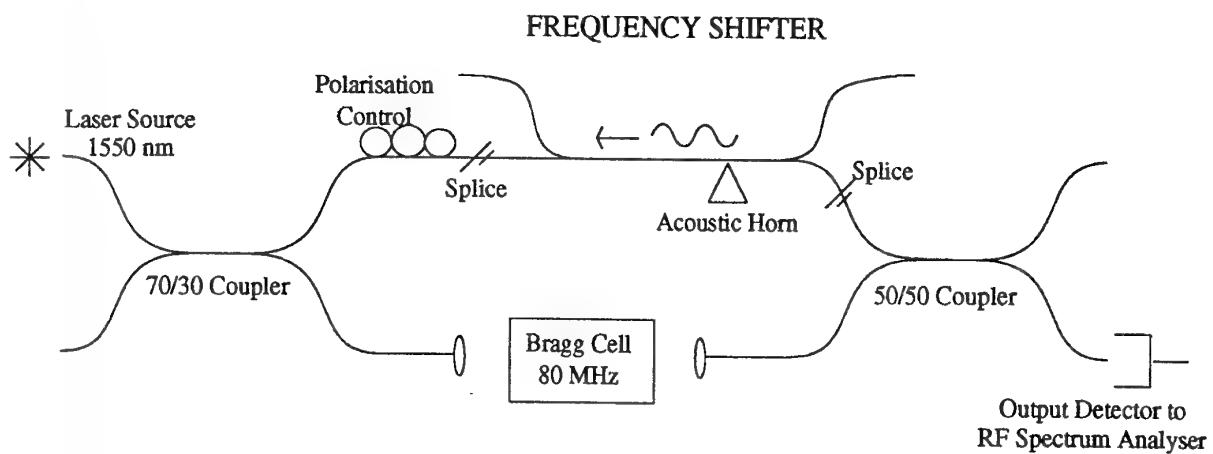


Fig. 49. Measurement of the frequency shift.

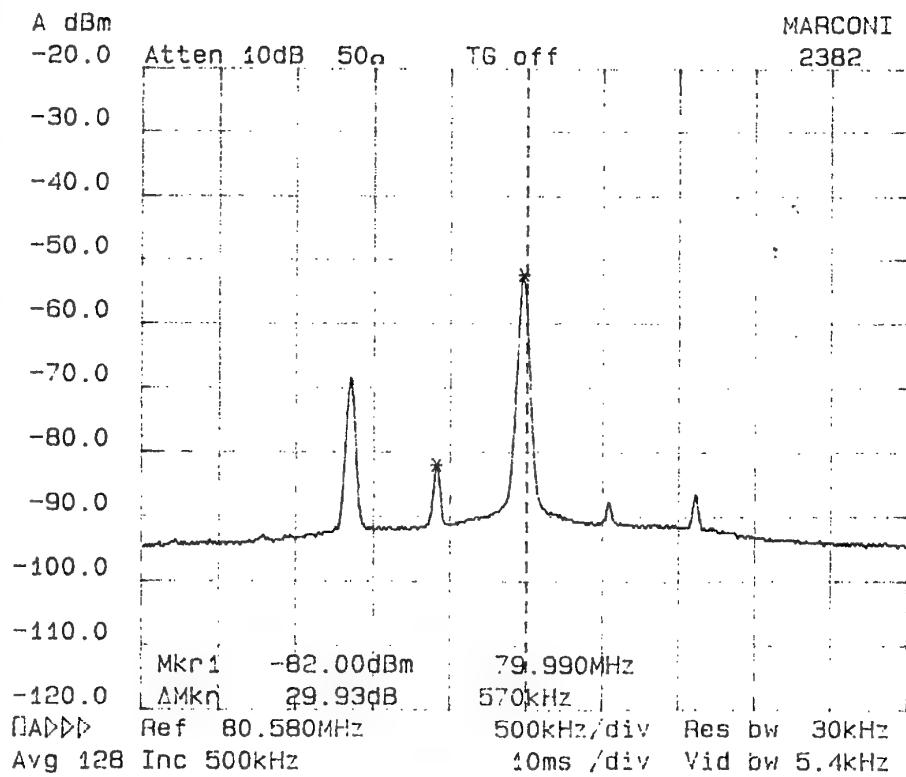


Fig. 50. RF spectrum of the output signal of the detector in the heterodyne experiment.

10.4 Frequency shifter at 12 MHz

The method of fabrication is identical to that outlined in previous sections - the only difference being the waist of the final device. For a resonant frequency of 10 MHz the waist diameter should be $5.5 \mu\text{m}$. The final device demonstrated a maximum splitting ratio of 1:16000 and an insertion loss of 0.03 dB. The diameter of the waist was finally measured by the SEM as $4.5 \mu\text{m}$, Fig. 51.



Fig. 51. SEM image of the coupler waist.

10.4.1 Acousto-optic resonance

The acousto-optic resonance was measured as described in section 10.3.1, the only difference being that the 1.25 MHz transducer was replaced by a 10 MHz PZT of diameter 8 mm. The base of the acoustic horn in this case measured 5 mm. The bond between horn and fibre was found to be more successful if the curing glue was replaced by cyano-acrylate adhesive. Maximum conversion was $\sim 99\%$ and was obtained for an acoustic frequency of 11.63 MHz. Fig. 52 is a graph of the performance of the acousto-optic frequency shifter as a function of pk-pk RF drive voltage across the PZT. The light was launched into the pretapered fibre.

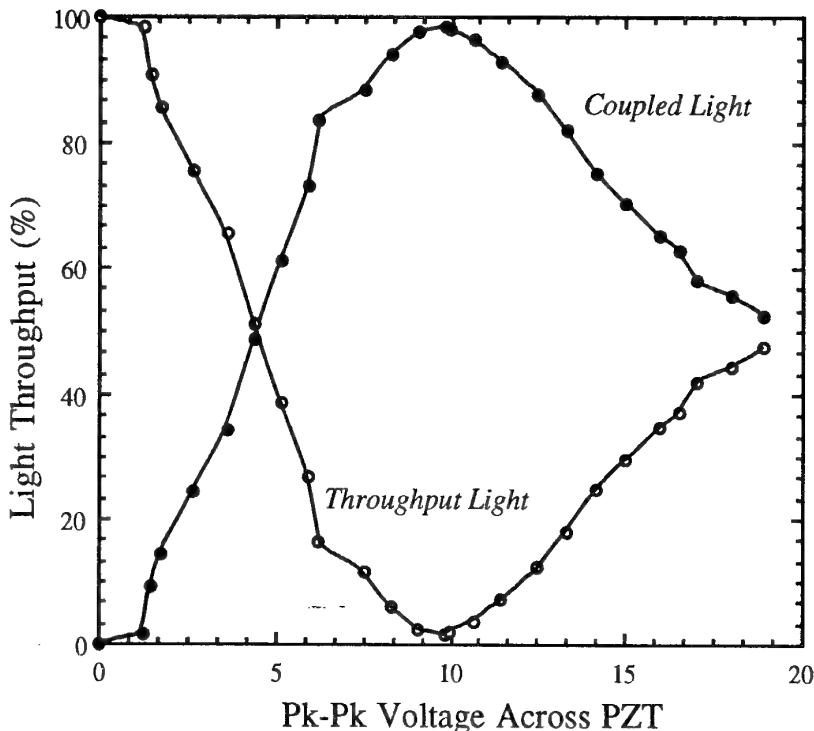


Fig. 52. Performance of the acousto-optic frequency shifter as a function of pk-pk drive voltage across the PZT.

10.4.2 Frequency shift

The 11.63 MHz frequency shifted signal was measured by incorporating the all-fibre Mach-Zehnder shown in Fig. 49. The beat signal was generated by mixing the 11.6 MHz frequency shifted signal with the output from the Bragg cell. The heterodyned signal was detected at one of the two output ports at the second coupler. The RF spectrum is shown in Fig. 53, with the upshifted signal at 68 MHz, and the carrier at 79.8 MHz. The carrier suppression was 40 dB. The image sideband at 91.6 MHz was suppressed by 30dB. However, there are unexpected sidebands at \sim 56 MHz and 118 MHz. At present these can only be explained by nonlinearities in the PZT transducer. In Fig. 54, the coupled port of the frequency shifter has been exchanged for the throughput port and the device turned off. The heterodyned signal at 79.5 MHz is clearly identified.

10.5 Conclusion

We have reported the successful construction of a null coupler using standard telecom fibre. Such a null coupler can be used to form an effective four-port acousto-optic frequency shifter. The procedure involves pretapering one of the identical pair of fibres to provide the necessary dissimilarity. In this way the device can be constructed entirely from one fibre type with four identical ports, and be compatible with any existing network. Our present device used standard 125 μ m single-mode telecom fibre throughout, for operation at a wavelength 1.55 μ m.

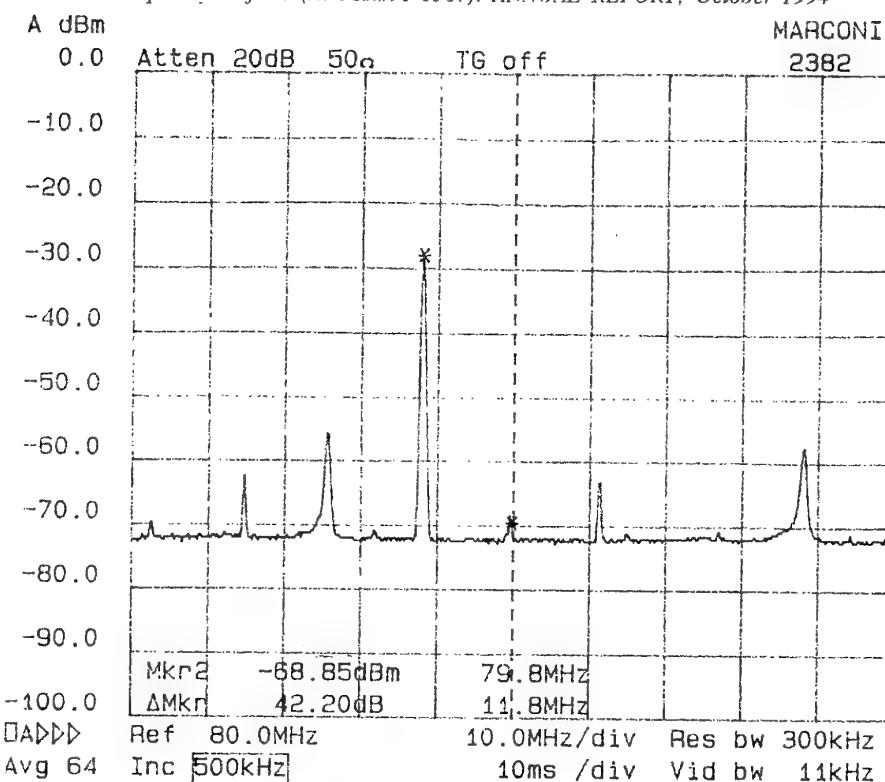


Fig. 53. RF spectrum of the output signal from the detector in the heterodyne experiment, at the coupled output of the frequency shifter.

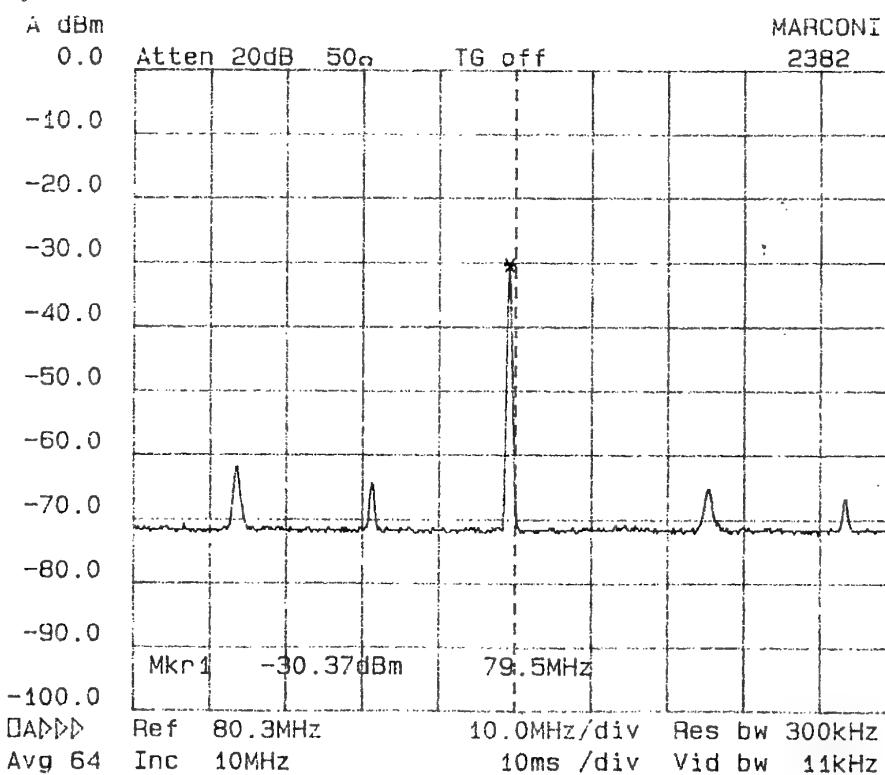


Fig. 54. RF spectrum of the output signal from the detector in the heterodyne experiment, at the throughput output of the frequency shifter with the device turned off.

11 POLARISATION PROPERTIES OF THE SECOND MODES

A short study was undertaken to investigate the polarisation properties of null couplers. These are expected to be significant, for reasons discussed in section 7.4. In that section, the coupler waist is assumed to support LP (scalar) modes. However, we will show here that for "narrow" circular coupler waists, the true hybrid (vector) modes must be considered.

11.1 "Wide" coupler waists

There are four degrees of freedom in the input and output fibres; that is, four independent ways in which light can be carried in the pair of fibres. These are the x and y polarisations of the fundamental mode of the un-pretapered fibre, and the x and y polarisations of the fundamental mode of the pretapered fibre. Passing into the coupler, the two fibres reduce in size and become more closely fused until the coupler waist is reached. The local modes change as the waveguide changes, but each local mode at every point along the coupler can usually be unambiguously identified with a local mode at the neighbouring point. The local mode is said to "evolve" along the transition. In general, the light couples between the local modes, as described by local mode coupling theory. However, if the transition is strictly adiabatic, then the light in each mode of the fibre remains in the same local mode as it evolves (this is the definition of "adiabatic").

Mode evolution in a null coupler is depicted schematically in Fig. 55. The evolution of the elliptical LP modes into circular hybrid modes is shown. The modes of the un-pretapered fibre evolve into the two fundamental modes of the waist, and the modes of the pretapered fibre evolve into the second modes of the waist. These are of course cladding modes; the fibre cores are too small to guide light, though they can still be significant enough to perturb the field distributions and propagation constants of the modes.

Where the waist is still elliptical, or the residual fibre cores are still significant (a "wide" waist), the lack of rotation symmetry imposes itself upon the character of the second modes. They are LP modes, one pair with lobes lying parallel to the plane of the coupler (call them LP_{11}' modes), and the other pair with lobes lying perpendicular to the plane of the coupler (call them LP_{11} modes). Each exists in two orthogonal plane-polarised states. In this case, excitation of the second modes is unambiguous. From excitation symmetry in the plane of the coupler, the LP_{11}' modes are never excited. If they ever become excited, any light remaining in them downstream of the coupler will be lost, Fig. 56.

Modal excitation can be determined by cleaving the coupler at the waist, and examining the far field pattern emerging from the cleaved end under different conditions. For a coupler with a "wide" waist supporting LP modes, input of light into the pretapered fibre should excite the LP_{11} modes, with lobes perpendicular to the coupler. Input of light into the unpretapered should excite the fundamental LP_{01} modes. This behaviour is illustrated in the photos of Fig. 45; this coupler had a "wide" waist.

The behaviour of a null coupler as it is twisted gives additional information about the second

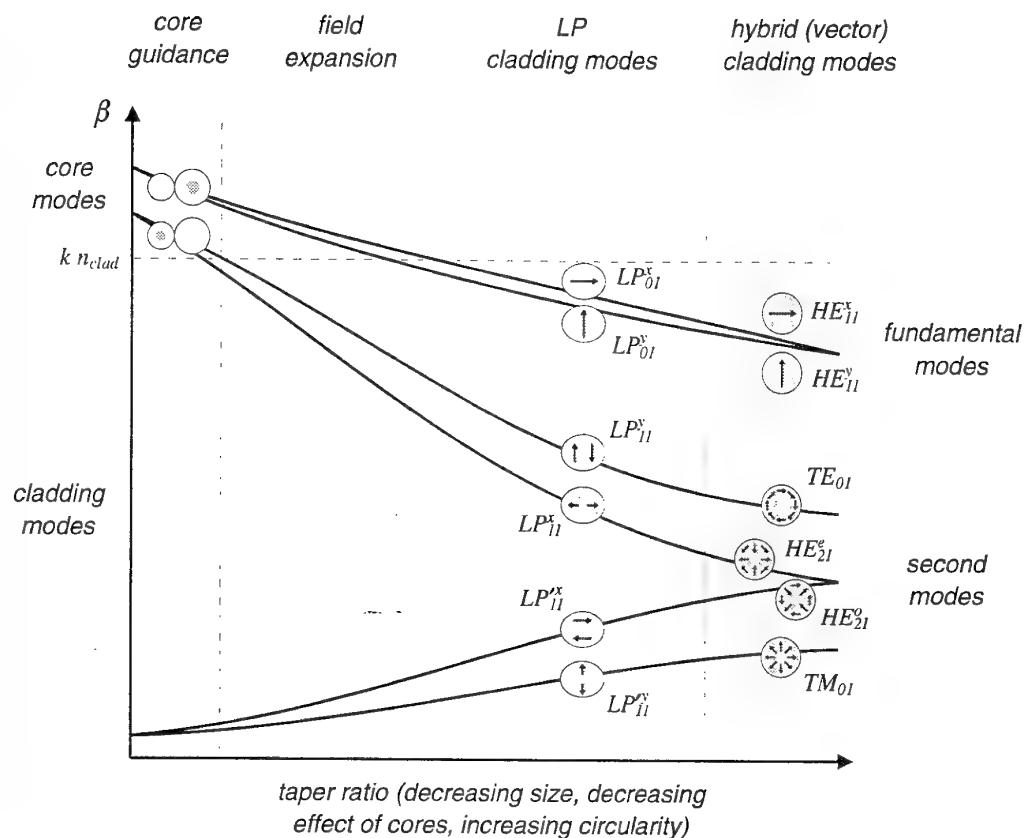


Fig. 55. The (schematic) progression of modes in a coupler from the unstretched fibres to a circularly symmetric narrow waist.

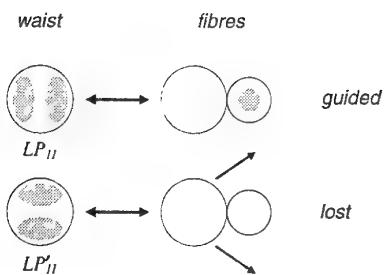


Fig. 56. The fate of light guided in the second mode of a "wide" (non-circular) coupler waist.

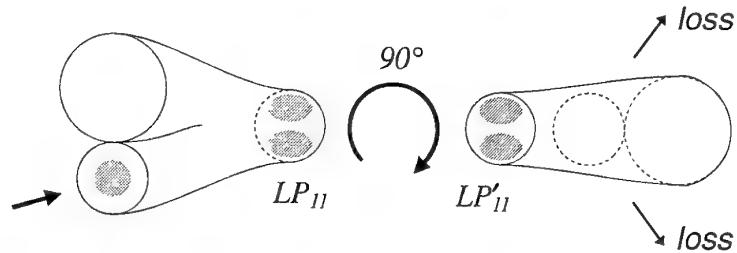


Fig. 57. The effect of twist on a circular LP_{II} mode.

modes if light is launched into the pretapered fibre. If we imagine the *LP* mode description to remain valid even if the waist is circular, then light in the LP_{11} mode excited in the start of the waist does not notice that the coupler is twisted. When the light reaches the end of the waist, the field pattern may no longer lie with its lobes perpendicular to the local plane of the coupler. If the coupler is twisted through 90° as shown in Fig. 57 the light is now in the LP_{11}' mode, and is therefore lost. We should see all input light being lost for twists at odd multiples of 90° , and all the light recovered at an even multiple of 90° . Another way to look at this is to say twist causes local mode coupling between the LP_{11} and LP_{11}' local modes. Since the modes are degenerate, there is no resistance to local mode coupling. Again, state of polarisation makes no contribution to this behaviour. Of course, the *LP* description is not valid for a truly circularly-symmetric waist, and in the non-circular waist the LP_{11} and LP_{11}' modes are not degenerate. This means there is some resistance to local mode coupling, and it takes more than 90° twist to reach a loss maximum, and that loss maximum will be less than 100%. Successive loss maxima will be greater as we move into a "strong twist" regime which overcomes the non-degeneracy of the modes. Nevertheless, it is true to say that at least 90° of twist are necessary to reach a loss maximum. The lost power can be expressed as:

$$P_{lost}(T) = \frac{4 I^2 T^2}{4 I^2 T^2 + l^2} \sin [4 I^2 T^2 + l^2]^{1/2} \pi \quad (63)$$

where T is the total number of twist revolutions imposed, I is an overlap integral and has its maximum value of unity for a circular cross section, and l is the length of the waist in multiples of the beat length between LP_{11} and LP_{11}' modes. In the inapplicable circular case, $l = 0$ since the beat length is infinite, and 100% loss occurs for $T = 1/4$ (i.e. 90°).

11.2 "Narrow" coupler waists

If the waist is circular in cross-section and the residual cores are truly negligible (a "narrow" waist) the set of second modes become hybrid modes and the *LP* description fails. This limit is shown on the right hand side of Fig. 55. The second modes which correspond to light guided in the pre-tapered fibre are no longer so distinct from those which are lost, and it is not obvious how they are excited by the input wave. Two (relatively) simple possibilities are (1) that the input light propagates adiabatically into the above *LP* modes where the coupler cross-section is still elliptical, but the hybrid modes are then excited as at an abrupt transition, or (2) that the light propagation is strictly adiabatic throughout the coupler. For both, the fundamental modes behave just as in the case of the "wide" elliptical waist - just the second modes are different.

11.2.1 Abrupt model

The abrupt model is depicted in Fig. 58. The combinations of hybrid modes which are excited by each *LP* pattern can be found by simple addition of the vector field at each point. We can see that unpolarised light in the pretapered fibre, which excites both LP_{11} modes (but neither LP_{11}' mode) will cause all four hybrid second modes to be excited in the coupler waist. Given the propagation constants at the right hand side of Fig. 55, there are three distinct beat lengths between excited second modes and the fundamental mode. Hence there will be three acousto-optic resonances, the central one being two-fold. Thus the abrupt model gives the experimentally observed triple resonance for the prototype frequency shifter, Fig. 40; a slight

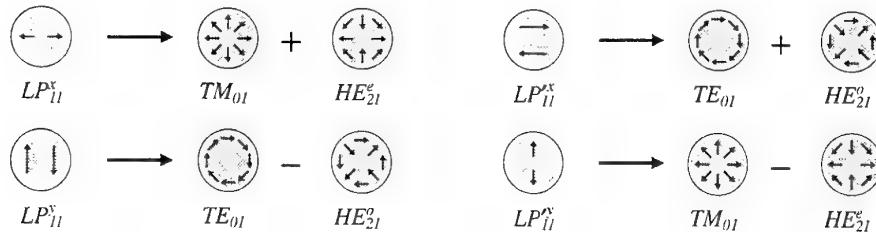


Fig. 58. The excitation of hybrid modes at the coupler waist, according to the abrupt model.

non-circularity would allow the central peak to split, giving the quadruple resonance sometimes observed. Unfortunately, the abrupt-model also gives loss. For example, suppose the input wave is in the LP_{11}^x mode. Because the TM_{01} and HE_{21}^e modes excited by it in the waist are not degenerate, they acquire a phase lag as they propagate along the waist. Unless they happen to be in phase again at the end of the waist, some of the output wave will be in the LP_{11}^y mode, which is lost. If they end up 180° out of phase, all of the output wave will be lost (Fig. 58). Hence we would see the loss of the coupler varying sinusoidally between 0% and 100% as it is being elongated on the rig (the calculated beat length between the hybrid modes is 14mm for the prototype coupler, so at least two or so beats would have been seen). This loss is an inevitable consequence of the non-adiabacity. We start with four degrees of freedom (two polarisations in each input fibre), expand that into six degrees of freedom (two fundamental and four second modes in the waist), and then try to compress it back into four (in the output fibres). This cannot be achieved except in special cases (for example when there happen to be a whole number of beat lengths along the waist).

11.2.2 Strictly adiabatic model

The strictly adiabatic model is depicted in Fig. 55 by simply following the evolution of the local modes from left to right, and assuming the light in a given local mode at the left hand side remains in that mode through to the right hand side. This is summarised in Fig. 59.

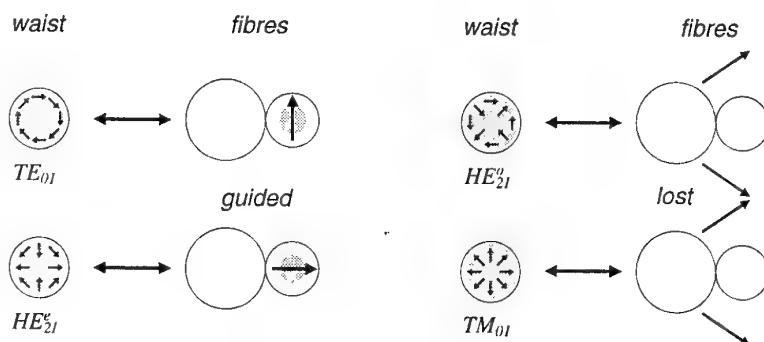


Fig. 59. The relation between hybrid modes in the waist and light in the fibres, for an adiabatic model.

Only two hybrid modes are excited, so we can no longer explain the triple acousto-optic resonance. But we can identify experiments on the passive null coupler to see if this model succeeds otherwise.

If the coupler waist is cleaved, there are a number of possible far field patterns which can be seen. For light entering in the un-pretapered fibre, the usual single spot is expected. For light entering in the pretapered fibre, some possible patterns are shown in Fig. 60; intermediate forms are also possible. The particular combination of hybrid modes excited in the coupler waist can be controlled by adjusting the polarisation of the light entering the pre-tapered fibre, as shown in Fig. 59. Note that if ever an *LP*-like two lobed pattern is seen, the lobes will be oriented at $\pm 45^\circ$ to the plane of the coupler. This contrasts with the case for the wide elliptical waist, where the lobes are always aligned perpendicular to the plane. Examination of the patterns through a polariser is also distinctive - if lobes are visible, the light is always polarised with the electric vector parallel to any lobes. This contrasts with a wide waist, where the light can be polarised in any direction with respect to the lobes. If just one hybrid mode is excited, it is possible to tell which with the help of the polariser. The doughnut pattern becomes a two-lobed pattern after passing through the polariser. If the polariser is then rotated, the two-lobed pattern rotates with the polariser for the TE_{01} mode, but against the polariser for the HE_{21}^e mode.

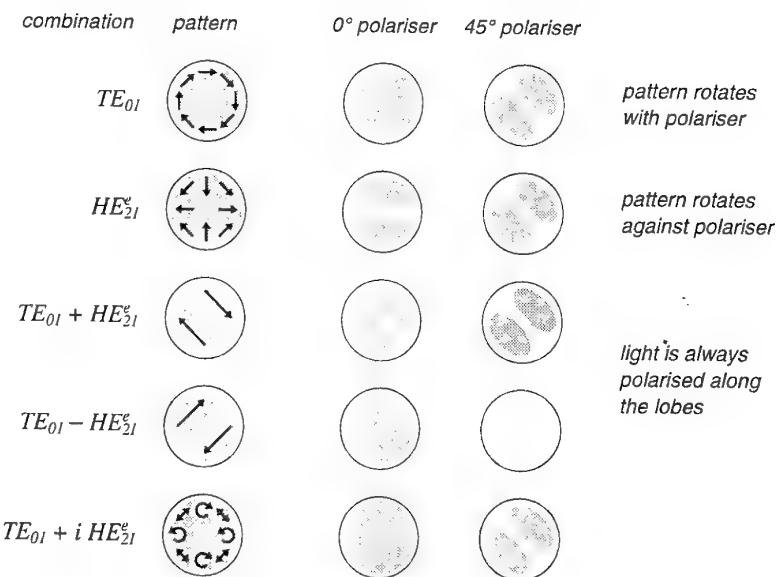


Fig. 60. Some possible far field patterns from a cleaved coupler waist, in the adiabatic model.

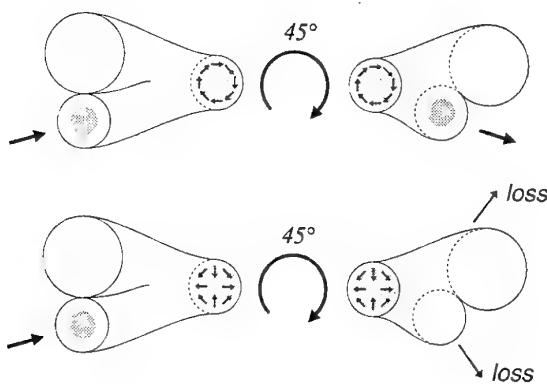


Fig. 61. The effect of twist on the excited hybrid second modes.

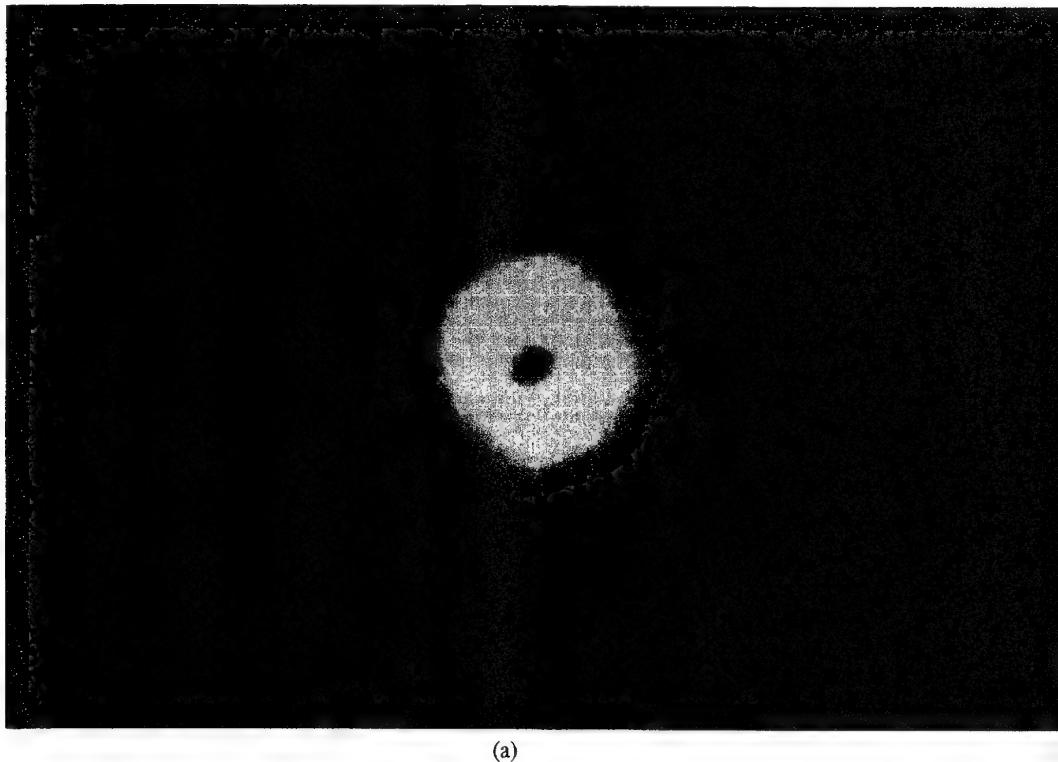
If the coupler is twisted, some more distinctive properties are predicted, Fig. 61. When just the TE_{01} mode is excited, the twist can have no geometrical effect on the relative orientation of the mode and the output fibres - this second mode has complete rotation symmetry. Hence there is never loss. However, the situation is different for the other possible second mode, the HE_{21}° mode. From the point of view of the local plane of the fibres, a 45° twist converts this mode to its degenerate partner, the $HE_{21}^{-\circ}$ mode. This mode corresponds to loss. Hence the model predicts that twist can cause loss, the maximum loss occurring at twist angles which are odd multiples of 45° (corresponding to $I = 2$ in the local mode coupler, equation (63)). The maximum amount of loss depends on the input polarisation. For x polarised light, this maximum loss is 100%. Note that the model with a wide elliptical waist can never give maximum loss for twist angles less than 90° , and the abrupt model for the narrow circular waist can never give more than 50% twist-induced loss.

11.3 Experimental

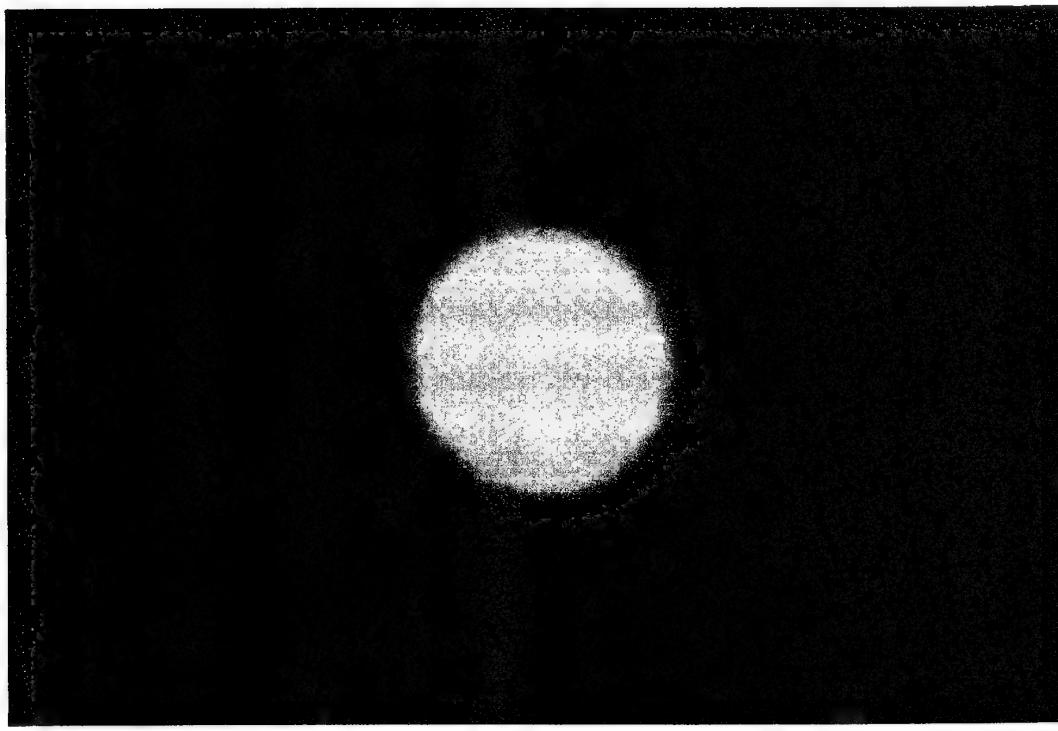
11.3.1 Observation of hybrid modes

A passive null coupler was made using fibre ND132_09 at $\lambda = 633$ nm as described in section 10.1. One fibre was stretched from $75 \mu\text{m}$ to $56 \mu\text{m}$, and the waist diameter of the final device was designed to be $\sim 2 \mu\text{m}$. The length of the final device was 25 mm, with an overall loss of 0.2 dB and a maximum splitting ratio of 1:4000. This coupler waist was considerably narrower than that of the couplers used in the frequency shifters.

After a series of twist experiments (described below) the coupler waist was cleaved and light emerging from the cleaved waist examined in the far field by projection on to a screen. When light was launched into the pretapered fibre, a doughnut far field pattern resembling a hybrid mode was seen, Fig. 62. For light launched into the un-pretapered fibre, the usual single spot was observed. Other second mode combinations are also possible; by adjusting the input polarisation to the null coupler it was possible to obtain two-lobed patterns oriented only at $\pm 45^\circ$ to the plane of the coupler, as expected for hybrid modes, Fig. 63. Second mode patterns

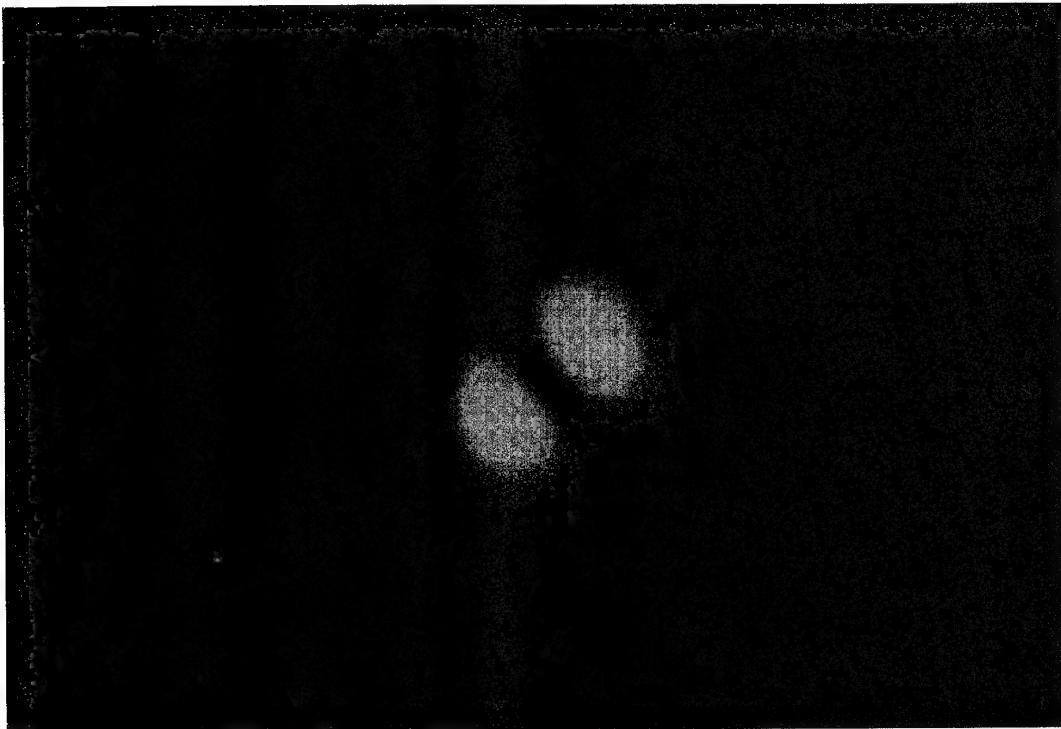


(a)

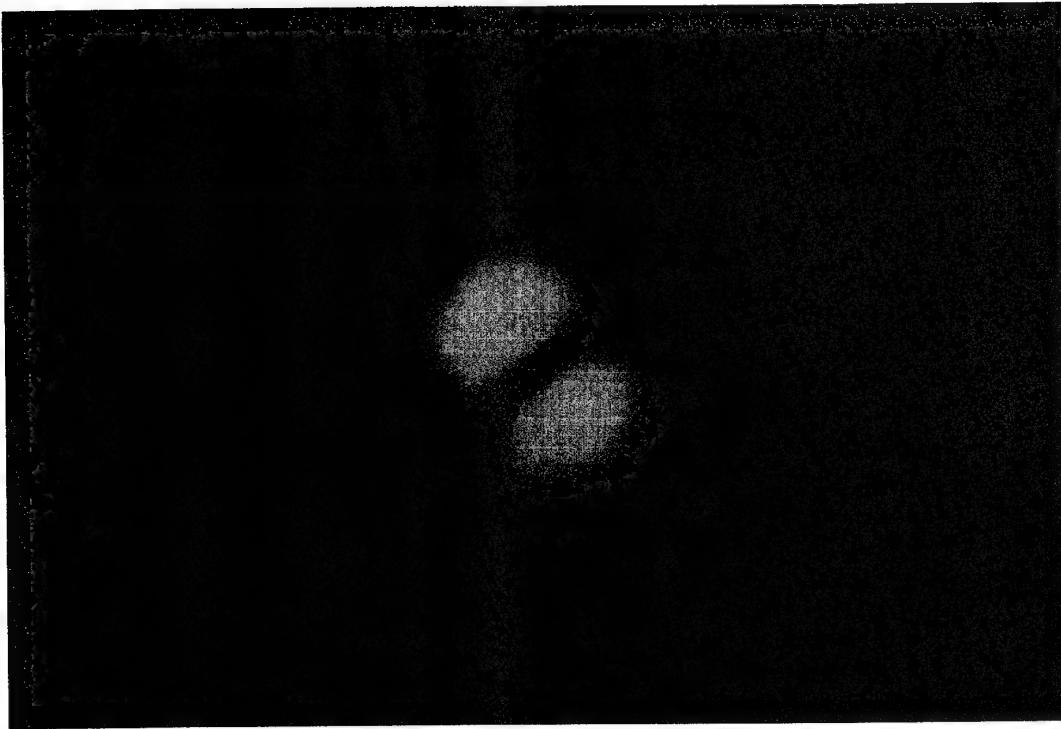


(b)

Fig. 62. Far field patterns emerging from the cleaved narrow waist of a null coupler a) for light launched into the pretapered fibre and b) for light launched into the un-pretapered fibre.



(a)



(b)

Fig. 63. Far field lobed like patterns at $\pm 45^\circ$ to the coupler axis corresponding to: a) $TE_{01} + HE_{21}^e$ and b) $TE_{01} - HE_{21}^e$. The light is launched into the pretapered fibre.

with lobes parallel or perpendicular to the axis of the coupler were not observed with this coupler.

Additional confirmation of the adiabatic model was obtained by inserting a polariser between the cleaved face and the screen, again confirming the behaviour anticipated in Fig. 60. With the input polarisation to the waist of the coupler adjusted to give a doughnut or ring-shaped pattern in the far field, inserting a polariser with its electric field vector perpendicular to the plane of the coupler can give far field patterns with a two lobed pattern perpendicular to the axis of the coupler, indicating the TE_{01} mode, or a two-lobed pattern lying parallel to the plane of the coupler, indicating the HE_{21}^e mode. Further information is obtained by rotating the polariser. In the case of the TE_{01} mode, the pattern rotates in the direction of rotation of the polariser, and for the HE_{21}^e mode the pattern rotates in the opposite sense.

In the case where the lobes are lying $\pm 45^\circ$ to the plane of the coupler without the polariser, (as shown in Fig. 63), when a polariser is inserted the light was found to be always polarised along the lobes. All these results indicate that this coupler waist is "narrow" as in the sense of section 11.2, and that the transition is completely adiabatic.

Fig. 64 shows the cleaved waist when examined by the SEM. The $2 \mu\text{m}$ waist is completely circular, and the fibre cores have become truly negligible.



Fig. 64. SEM of the $2 \mu\text{m}$ cleaved waist.

11.3.2 The effect of twist on the null coupler

We expect the output transmission behaviour of the coupler (for input light in the pretapered fibre) to determine whether the coupler waist is "wide" or "narrow". Light is passed through a polarisation controller before entering the coupler. One end of the null coupler is supported by a rotation stage, the opposite end being fixed. For the coupler with the "narrow" waist described above, with an arbitrary input polarisation state, it was found that loss varied sinusoidally with twist angle, with maximum loss for a twist angle of $\sim 45^\circ$. This is the behaviour expected for a narrow waist supporting hybrid modes. With the twist set for maximum loss, the polarisation controller was adjusted to give the minimum throughput (close to zero), indicating that the light at the waist is in a near pure HE_{21} mode. The graph of throughput versus twist for this polarisation state is shown in Fig. 65.

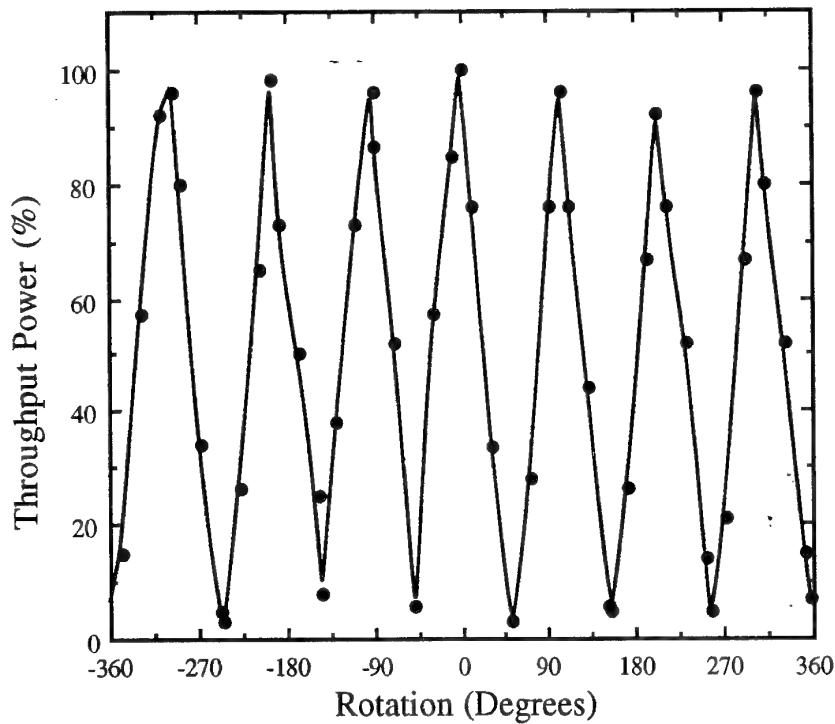


Fig. 65. Response of a null coupler with twist, with the HE_{21} mode excited.

A similar set of results can also be taken when the TE_{01} mode is excited. This is achieved by adjusting a half wave plate at the laser source to give the orthogonal polarisation, and hence a nearly pure TE_{01} mode in the waist. The throughput versus twist is shown in Fig. 66. From Fig. 66 the TE_{01} mode is seen to experience minimal loss and confirms the strong rotational symmetry anticipated.

In the case of a "wide" elliptical waist, it is understood that local mode coupling takes place between the LP_{11} and LP_{11}' modes. For a null coupler with a waist of diameter $9.5 \mu\text{m}$, the throughput against twist is shown in Fig. 67. This coupler was made with fibre ND 132_09,

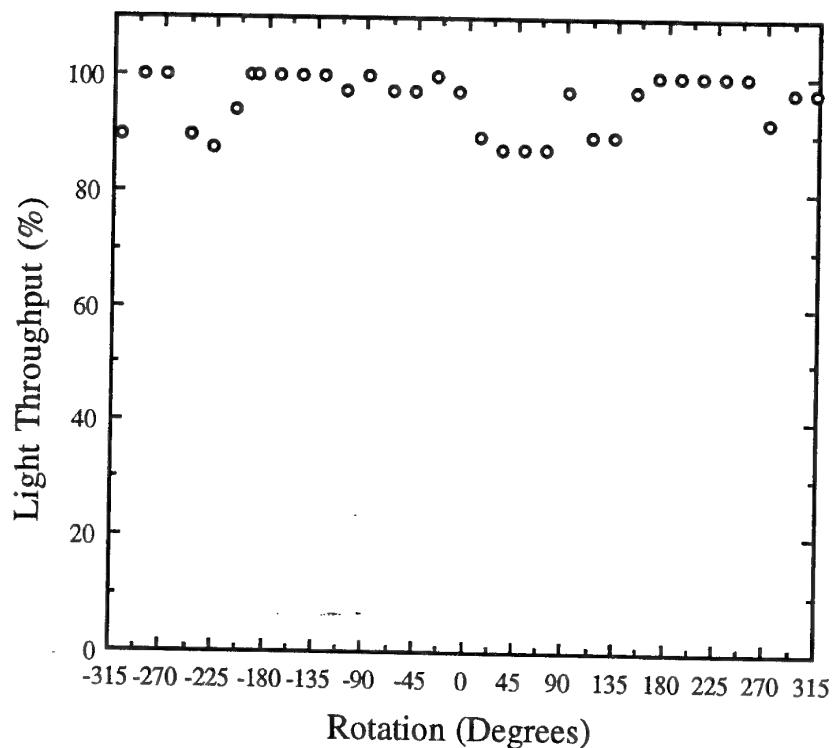


Fig. 66. Response of the null coupler with twist, with the TE_{01} mode excited.

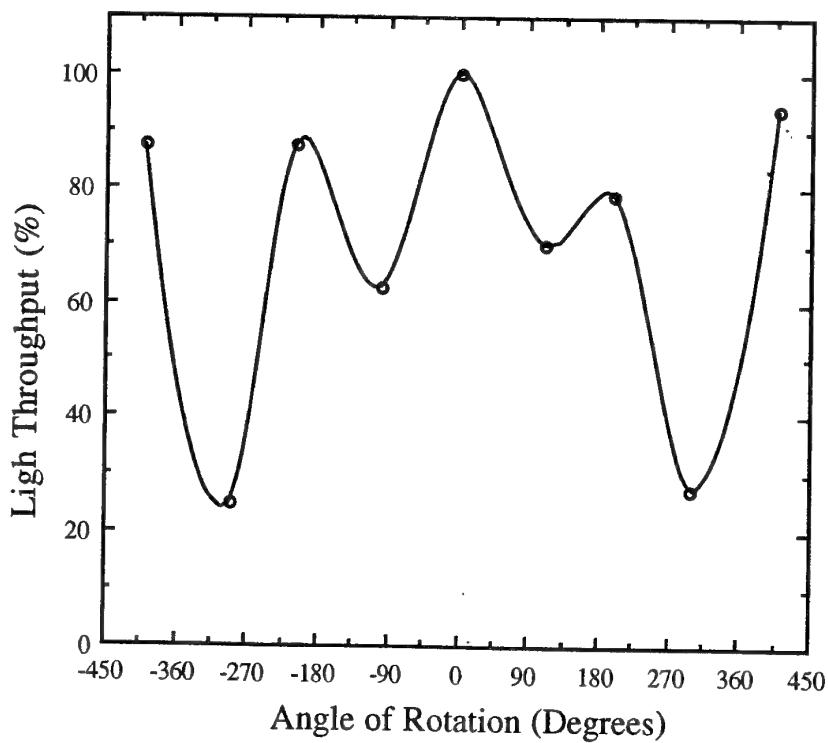


Fig. 67. Response of a "wide" null coupler with twist.

and had an excess loss of 0.3 dB and a maximum splitting ratio of 1:700. It can be seen that the first minimum occurred at twist angles marginally greater than 90°. As the twist is increased, successive minima become stronger because strong twists overcome the non-degeneracy of the modes. This is exactly the behaviour indicated for a "wide" coupler waist.

11.4 Conclusion

Second modes in the waist of null couplers can show either *LP* mode properties or hybrid mode properties. These will have to be considered to fully understand the behaviour of null coupler frequency shifters. Although the triple-peaked spectra of Figs. 40, 47 & 48 clearly result from a polarisation effect, we are still unable to determine exactly how they arise.

12. CONCLUSIONS AND PROPOSED FUTURE WORK

12.1 Current state of progress

We now review the project goals outlined in section 2.1, and compare the specifications achieved with the prototype frequency shifter.

Low insertion loss: The insertion loss of the prototype, <0.2 dB, is low. This goal is already satisfied.

Carrier suppression >50 dB: Carrier and image sideband suppressions greater than 30 dB have been measured. Goal not yet satisfied.

Four- or three-port device: By its nature, the null coupler device automatically has four ports. Already satisfied.

Identical standard single-mode fibres: Although the prototype was made from a pair of dissimilar and non-standard fibres, subsequent devices described in section 10 have been demonstrated with identical ports made from standard telecoms fibre, for single-mode operation at the wavelength of 1550 nm. Already satisfied.

Low electrical drive power: The electrical drive power for the prototype was just 1 mW. Although much greater than the theoretical requirement of 170 nW, this is in turn much smaller than any rival design. Already satisfied.

Robust and stable package, no longer than 30 cm: The device is less than 30 cm long, but is unpackaged. Not yet satisfied.

Polarisation dependence: This was not specified.

High frequency operation: This was not specified.

Additional, unspecified features of the prototype device are:

- simple monolithic construction (no splices needed);
- usable as a switch, modulator or tunable filter, as well as a frequency shifter.

We propose to continue our research with the aim of realising these specifications in a deliverable device. The work required to satisfy each requirement is now outlined.

12.2 Carrier and image sideband suppression

The presence of unshifted light is due to deviation of the real coupler from null behaviour; that is, whether the coupler couples any light when passive. If the acousto-optic coupling between fundamental and second modes is not complete (because either the acoustic amplitude or the

polarisation control is not ideal), or if a third mode is excited, then some unshifted light can emerge with the frequency-shifted light. Deviation from null behaviour can be intrinsic to the coupler, or can be caused by inadequate mounting of the coupler (causing bends and/or twists) or mode-conversion due to surface scatterers (such as dust particles or microscopic cracks).

The best null couplers so far made had maximum coupling ratios of less than 1:6000. With such a coupler, an acousto-optic conversion efficiency of just 94 % would yield the target carrier suppression of 50 dB. For the 1:400 null coupler in the prototype device, a conversion efficiency of 99 % should have given a carrier suppression of 46 dB. The measured value was only 30 dB. We think this is due to the excitation of a third mode. The evidence for this was that the small amount of light coupled in the passive device appeared to be carried in the second mode of the two-moded fibre. If the device had been made from two single-mode fibres, this light would have been lost (contributing an insignificant additional 0.01 dB to the insertion loss) and so would not have contaminated the frequency-shifted output. Thus the target of 50 dB suppression should be readily achievable with a truly single-mode null coupler with the type of performance already achieved. The 1550 nm device of section 10.4 gave a carrier suppression of 40 dB.

Image sideband contamination can be caused by deviations from null behaviour as described above, but the major cause is bi-directionality of the flexural acoustic wave along the coupler waist. The mechanism is simple: if an acoustic wave travelling in one direction along the coupler waist causes light to couple with a frequency shift in one sense, then any acoustic wave in the other direction couples light with the opposite frequency shift. In our device, the acoustic wave generated by the transducer travels in only one direction along the coupler waist, but if some of the wave is reflected at the far side of the coupler, then there will also be a backward propagating acoustic wave in the coupler waist. This acoustic reflection will have to be reduced.

Acoustic reflection is due to abrupt changes in the acoustic velocity along the fibres, and an absence of acoustic losses. We consider that the greatest source of acoustic reflections are the points along the bare fibres where the polymer buffer coating resumes (the acoustic energy transmitted beyond these points will be rapidly dissipated in the coating). We will reduce acoustic reflections by making the transitions between bare and coated fibres more gradual, possibly by careful addition of coating material in a "tapered" fashion. Image sideband contamination due to deviations of the coupler from null behaviour are minimal; for the 1:400 coupler of the prototype, this source (taken alone) would already satisfy our suppression target of 50 dB.

12.3 Robust and stable package

This is a very important requirement if the device is to be developed from a laboratory experiment to a practical device. Instabilities in performance result from small mechanical deformations, which can change the nature of the coupler waist as an optical or acoustic waveguide. Any changes in the optical beatlength or the acoustic wavelength will in turn change the acousto-optic resonance condition, and detune the device. These mechanical deformations can be caused directly if external forces are applied to the package, or indirectly by the differential thermal expansion of the various parts of the package, the fixing glue and the silica fibres and coupler. The package therefore has to be stiff, of uniform construction, and of a material with similar thermal expansion properties to silica glass. These issues have already been addressed for conventional fused taper couplers, although for less critical tolerances, and

we can adapt existing coupler packaging technology for the frequency-shifter. Couplers are usually housed in a primary package made of the low expansion alloy invar, or of silica itself. This minimises thermal strains. The primary package is then held inside a steel secondary package for strength. As well as protecting from damage by bending or crushing, such a secondary package will also keep dust and other particles from the exposed surface of the coupler, preventing surface damage and the scattering of light.

However, there remains the risk of inertial damage, due to sharp knocking or dropping for example. Some degree of cushioning may be possible with rubber housings, but for a comprehensive approach to making our device rugged we may have to seek assistance from mechanical engineers more experienced in such problems. It is useful to bear in mind that close-channel WDM couplers, with waist diameters not much greater than that of the coupler in our device, have been satisfactorily packaged by various manufacturers and are in routine use.

The performance of the device can also be made less sensitive to the environment if it is made as broadband as possible. This in turn broadens the resonance condition, so small changes will have less effect. This is discussed further under the next section.

The sensitivity of the acoustic wavelength (and hence the resonance condition) to fibre tension can be used to advantage as a tuning mechanism. The effects of tension will be investigated quantitatively, and compared with (34).

12.4 Broadband operation

This design of frequency-shifter is always very broadband in one sense, that a given device can be operated over a very wide range of RF frequencies and optical wavelengths. However, if (say) the frequency is changed, the wavelength of operation must also change to compensate.

The null coupler is made using a technique designed to give the coupler a waist with a uniform diameter along its length. Such uniformity is needed to give a strong interaction, but in theory results in very narrow-band operation. In a perfectly uniform coupler waist with the dimensions of that in the prototype device, the interaction has an optical bandwidth of 3.5 nm and an acoustic bandwidth of 0.02 MHz about a given resonance.

These bandwidths can be increased if the coupler waist is not perfectly uniform. Indeed, the optical bandwidth of the real prototype was about 10 nm, which is due to slight non-uniformities (of the order of just 0.05 μm) introduced unintentionally by the fibre elongation process. Bandwidth can be further increased by deliberately making the coupler waist slightly narrower at one end than the other (Fig. 68). This would reduce the constraints on optical and RF sources, and make the device less sensitive to environmentally-caused changes to the resonance condition. A greater acoustic power would be needed to give the same strength of interaction for each wavelength over the shortened effective length, but this is not a problem because the device has such a low power requirement anyway.

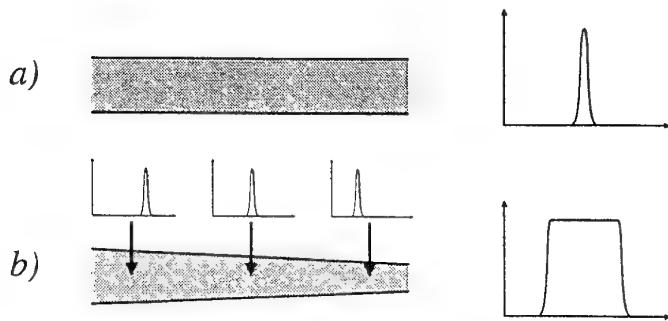


Fig. 68. The broadening of an acousto-optic resonance, by making the taper waist non-uniform (b) instead of uniform (a). Different parts of the non-uniform waist contribute a resonance at a different wavelength. The x axes of the graphs could represent either optical wavelength for a given acoustic frequency, or vice versa; the y axes represent coupled (and therefore frequency-shifted) power.

12.5 Polarisation dependence

The acousto-optic interaction in a null taper coupler is unavoidably polarisation-dependent. The resonance condition is in fact split into two resonances, one for each eigen-polarisation (Fig. 69a). For our prototype, this results in the prediction of two resonances spaced by 18 nm in optical wavelength. (In practice, *three* resonances were observed. We have as yet only a partial explanation for this. It will be a minor yet important part of the project to conduct experiments designed purely to complete our understanding of this effect.)

Methods for reducing or eliminating this dependence have been proposed, but none have practical merit. We see three alternatives for coping with this. Firstly, the device can be used with polarised light, with polarisation controllers to match the system polarisation to one of the device's eigen-polarisations. Secondly, the polarisation dependence can be ignored in any system with effectively unpolarised light, provided it is acceptable that a maximum of just 50 % of the input light (ie, only one polarisation) is frequency-shifted. (The other, unshifted, polarisation will not contaminate the shifted wave; remember that unshifted light will emerge from the other output port of the device.)

Thirdly, it should be possible to mask the polarisation dependence in the same way that broadband operation can be achieved, with a non-uniform coupler waist. If the device's bandwidth exceeds the polarisation splitting, there will be an overlap where both polarisations are efficiently frequency-shifted; one polarisation is coupled within the wider end of the waist, the other being coupled within the narrower end (Fig. 69b). More directly, the coupler waist could comprise two successive uniform regions with different diameters, each designed to give acousto-optic resonance for a different eigen-polarisation, but for the same optical wavelength and acoustic frequency (Fig. 69c). Either way, the resulting frequency-shifter could behave in an effectively polarisation insensitive manner for restricted combinations of wavelength and frequency.

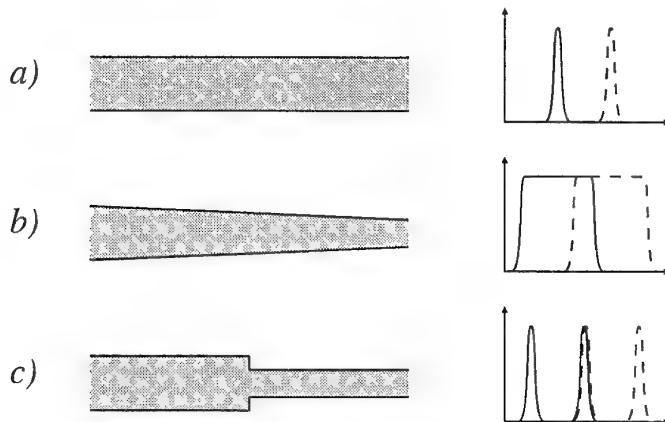


Fig. 69. Masking of polarisation dependence. A uniform coupler waist (a) gives two narrow resonances, one for each polarisation. In a non-uniform waist (b), each resonance is broadened. In a waist comprising two uniform regions (c), the resonance is two-peaked for each polarisation. The device is polarisation insensitive where the resonances overlap.

12.6 Low electrical drive power

The demonstrated drive power requirement of 1 mW is over 80 times smaller than the reported competition, and is probably low enough for practical purposes. However, the theoretical power requirement of the prototype is just 170 nW, and the actual acoustic power required was about $2 \mu\text{W}$, giving 0.2 % efficiency for the conversion of electrical power into power in the flexural acoustic wave at the coupler waist. The transducer arrangement in our prototype was not optimised in any way, whereas an electro-acoustic conversion efficiency of 11 % has been reported for Stanford's dual-mode fibre frequency shifter¹². There is therefore plenty of scope for reduction of the drive power, which would be worthwhile in itself and will give further room for manoeuvre for broadband operation and the masking of the polarisation dependence.

There is a long chain of elements between the RF source and the coupler waist, and improvements can be made in most of them.

50 Ω RF source to PZT disc: A simple electrical impedance matching network would minimise the drive voltage and losses in the transmission line. The network could be active, or maybe just a transformer and LC circuit. In any case, the techniques are well known, and will be applied.

The PZT disc: This can be chosen to have a thickness frequency matching the required drive frequency. These discs have an extremely frequency-dependent response, which will limit the RF bandwidth of the device, so techniques from the acoustics literature may be needed to smooth this frequency response. Some of these involve careful shaping of the mount upon which the base of the disc is fixed. Aside from such considerations, it is otherwise desirable to maximise the acoustic mis-match between disc and mount, in order to minimise the dissipation of acoustic power into the mount.

PZT disc to concentrator horn: These are presently fixed together using epoxy resin. Even a thin layer of this can significantly reduce transmission of acoustic energy. Alternative adhesives, methods for reducing the thickness of adhesive, and possibly adhesive-less

techniques such as indium bonding, can all be explored.

The concentrator horn: this is presently made of aluminium, which provides an adequate acoustic match to both the PZT ceramic and silica glass. It may be worthwhile to study refinements to the simple conical shape, to ensure that the maximum amount of energy entering the base is transmitted as vertical motion at the tip.

Concentrator horn to fibre pair: in the prototype, these were weakly fixed together by a relatively large mass of UV-curing glue. It is unlikely that there was direct vertical contact between the horn tip and either fibre. It is not clear exactly how this junction can be improved, but it is probably here that the greatest improvements in drive efficiency can be made.

Along the fibre pair: the propagation of a flexural acoustic wave along a pair of independent parallel fibres is not straightforward. However, there is little that can be done to change this link in the chain of power transfer. One possibility would be to provide a region at one end of the coupler where the two fibres are fused together but not stretched, and excite the acoustic wave there. This would however not be simple to achieve without degrading the coupler.

Through the coupler's taper transitions: the transitions might be expected to be a source of acoustic reflections and hence loss, but acoustic propagation here is actually very efficient. This was verified in experiments with a single-taper, where the measured acoustic amplitudes in the un-narrowed fibre and in the taper waist corresponded to virtually the same acoustic power.

12.7 High frequency operation

Operation at frequencies greater than a few MHz is not a formal goal of the project. However, the null coupler design can in principle be operated at frequencies approaching 1 GHz, which are much greater than is possible in competing all-fibre devices. It might therefore be appropriate to study high-frequency shifting. The main change to the existing design that would be required would be a change in the transducer; PZT discs are inappropriate for such frequencies.

13. ANCILLARY TOPICS

13.1 Publications

L. Dong, T. A. Birks, M. H. Ober and P. St. J. Russell: "*Intermodal coupling by periodic microbending in dual-core fibers - comparison of experiment and theory*", IEEE Journal of Lightwave Technology, **12**, 1994, pp. 24-27. This paper describes some of the theory behind the dual-core frequency shifter.

T. A. Birks, P. St. J. Russell and C. N. Pannell: "*Low power acousto-optic device in tapered single-mode fiber*", Proceedings of the Conference on Optical Fiber Communication (OFC '94), San Jose CA, February 1994, pp. 98-99. This paper describes the acousto-optic device based on

a single taper.

T. A. Birks, P. St. J. Russell and C. N. Pannell: "*Low power acousto-optic device based on a tapered single-mode fiber*", IEEE Photonics Technology Letters, **6**, 1994, pp. 725-727. This is an expanded version of the above OFC paper.

T. A. Birks, S. G. Farwell, P. St. J. Russell and C. N. Pannell: "*Fibre acousto-optic frequency shifter based on a null coupler*", Proceedings of the Optical Fibre Sensors Conference (OFS10), Glasgow UK, October 1994, pp. 281-284. This paper describes the prototype null coupler frequency shifter.

T. A. Birks, S. G. Farwell, P. St. J. Russell and C. N. Pannell: "*Four-port fiber frequency-shifter with a null taper coupler*", due to appear in Optics Letters, **19**, 1994. This is an expanded version of the above OFS paper.

D. O. Culverhouse, S. G. Farwell, T. A. Birks, P. St.J. Russell and C. N. Pannell: "*High performance 4-port acousto-optic frequency shifter at 1.55 microns using standard single-mode telecom fibre*", submitted to the Conference on Lasers and Electro-Optics (CLEO '95), Baltimore, to be held in May 1995.

We also intend to write a paper summarising the theory of the null coupler frequency shifter, once we have reached a thorough understanding of its polarisation properties.

13.2 Participants

Dr Tim Birks was the research fellow working full-time on the project with the principal investigator, Dr Philip Russell. His involvement has been scaled down, following the arrival of Dr David Culverhouse in June 1994 as full-time research fellow. Selina Farwell joined the project as a Ph.D. student in November 1993. Within the ORC, Dr Chris Pannell was closely involved in the project, though he has recently left to take up a readership at the University of Kent. The dual-core fibres were fabricated by Dr Liang Dong. There have been no interactions with other groups.

13.3 Discoveries, inventions, patents

We believe the acousto-optic device based on the null coupler to be significant as a switch, modulator and tunable filter as well as a four-port frequency shifter. It also appears to be novel. We have therefore prepared a wide-ranging patent on the device, which was filed in the U.K. on the 18th February 1994 under the title "Acousto-optic device".

13.4 References

1. B. Y. Kim, J. N. Blake, H. E. Engan and H. J. Shaw: "*All-fiber acousto-optic frequency shifter*", Optics Letters, **11**, 1986, pp. 389-391.
2. J. Blake and P. Siemsen: "*Practical compact high performance fiber-optic frequency shifter*", Proceedings of the Optical Fiber Sensors Conference (OFS9), Florence Italy, 1993, pp. 301-304.
3. H. Sabert, L. Dong and P. St. J. Russell: "*Versatile acousto-optical flexural wave-modulator, filter and frequency shifter in dual-core fibre*", International Journal of Optoelectronics, **7**, 1992, pp. 189-194.
4. A. W. Snyder and X.-H. Zheng: "*Fused couplers of arbitrary cross-section*", Electronics Letters, **21**, 1985, pp. 1079-1080.
5. J. N. Blake, B. Y. Kim, H. E. Engan and H. J. Shaw: "*Analysis of intermodal coupling in a two-mode fiber with periodic microbends*", Optics Letters, **12**, 1987, pp. 281-283.
6. H. F. Taylor: "*Bending effects in optical fibers*", IEEE Journal of Lightwave Technology, **2**, 1984, pp. 617-628.
7. W. K. Burns and A. F. Milton: "*Mode conversion in planar-dielectric separating waveguides*", IEEE Journal of Quantum Electronics, **11**, 1975, pp. 32-39.
8. T. A. Birks and Y. W. Li: "*The shape of fiber tapers*", IEEE Journal of Lightwave Technology, **10**, 1992, pp. 432-438.
9. J. D. Love: "*Application of a low-loss criterion to optical waveguides and devices*", IEE Proceedings Part J, **4**, 1989, pp. 225-228.
10. A. W. Snyder and X. H. Zheng: "*Optical fibers of arbitrary cross sections*", Journal of the Optical Society of America, **3**, 1986, pp. 600-609.
11. A. W. Snyder and J. D. Love: "*Optical waveguide theory*", Chapman and Hall, London, 1983.
12. J. N. Blake, B. Y. Kim, H. E. Engan, H. J. Shaw: "*Analysis of intermodal coupling in a two-mode fiber with periodic microbends*", Optics Letters, **12**, 1987, pp. 281-283.
13. D. B. Mortimore: "*Wavelength-flattened fused couplers*", Electronics Letters, **21**, 1985, pp. 742-743.

14. APPENDIX: TABLE OF SYMBOLS

A_∞	the transverse plane in a fibre
a	radius of a fibre core
a_1, a_2	radii of the cores of a dual-core fibre
α	rate at which hot-zone length increases during second stage of taper elongation
β	propagation constant
β_e	propagation constant of the even (fundamental) mode
β_o	propagation constant of the odd (second) mode
β_1, β_2	propagation constants of the modes of the cores of a dual-core fibre
$\bar{\beta}$	"average propagation constant" of Bloch modes
β_\pm	propagation constants of Bloch modes
C	acousto-optic coupling coefficient
c_{ext}	speed of extensional waves, = 5760 m s ⁻¹ for silica
χ	an elasto-optic constant, = -0.22 for silica
d	core separation in a dual-core fibre
$\Delta\beta_x, \Delta\beta_y$	$\beta_e - \beta_o$ for light in a coupler polarised parallel to the x and y axes
$\Delta f, \Delta\lambda, \Delta(2R), \Delta\Lambda$	specific cases of Δq
Δn	change in refractive index
$\Delta n(x, y)$	peak change of refractive index due to an acoustic wave
$\Delta n(x, y, z, t)$	change in refractive index due to an acoustic wave
$\Delta\nu$	optical frequency shift
δ	acousto-optic mismatch (deviation from resonance)
Δq	a range of q , bandwidth or polarisation splitting (depending on the context)
E	amplitude of a light wave in the fibre vibrometer
ε	a phase function
η	a phase constant in the fibre vibrometer
f	acoustic frequency or RF drive frequency
f_{max}	theoretical maximum frequency shift
I	local-mode-coupling overlap integral
I_O	acousto-optic overlap integral
I_0, I_1	modified Bessel functions of the first kind
i	$\sqrt{(-1)}$
J_0, J_1	Bessel functions of the first kind

j_{01}, j_{11}	first zeros of Bessel functions J_0 and J_1 respectively
K_0, K_1	modified Bessel functions of the second kind
k	free space optical wave constant, $= 2\pi/\lambda$
κ	acoustic wave constant, $= 2\pi/\Lambda$
L	interaction length, length of taper waist
L_B	beat length of even and odd modes, $= 2\pi/(\beta_e - \beta_o)$
L_o	hot-zone length at the start of taper elongation
$L(x)$	variation of hot zone length during taper elongation
l	length of a coupler waist, in multiples of the LP_{11} and LP'_{11} beatlength
Λ	wavelength of acoustic wave
λ	optical wavelength <i>in vacuo</i>
λ_{cutoff}	second mode cut-off wavelength of a single-mode fibre core
M	maximum intrinsic coupling
m_q	power to which q is raised in the resonance condition (8), $m_f \equiv 1$
N	an integer
NA	numerical aperture, $= \sqrt{n_{core}^2 - n_{cladding}^2}$
n_{air}	refractive index of air, $= 1$
$n_{core}, n_{cladding}$	core and cladding refractive indices in a fibre
n_o	refractive index of silica, $= 1.458$
ν	optical frequency
Ω	acoustic angular frequency
ω	optical angular frequency
P	power carried by acoustic wave
P_t, P_c	normalised optical power in the throughput and coupled outputs of a coupler
$P_{in}, P_{fs}, P_{carrier}$	input, shifted output, and unshifted output powers of a frequency shifter
$P_j(\Delta\nu)$	output power in port $j = 1, 2$ with frequency shift $\Delta\nu = 0, +f, -f$
P_o	output optical power of the fibre vibrometer
Φ	phase difference acquired by two modes propagating along a coupler waist
ϕ	Bloch mode excitation parameter
φ	azimuthal angle in polar co-ordinates
Ψ_{\pm}	normalised field distribution functions of Bloch modes
$\psi(z)$	field distribution function at z
ψ_e, ψ_o	normalised field distribution functions of the even and odd normal modes
ψ_1, ψ_2	normalised field distribution functions of light entering/leaving ports 1 and 2
q	generic parameter affecting acousto-optic resonance (λ, f, R, n_o or Λ)
R	radius of taper waist or electrical resistance (depending on the context)
R_{ac}	acoustic power reflection coefficient

<i>r</i>	radial distance in polar co-ordinates
<i>r</i> _o	initial radius of a fibre being tapered singly
<i>r</i> ₁ , <i>r</i> ₂	initial radii of fibres, or optical amplitude reflection coefficients in vibrometer
<i>ρ</i>	density of silica, = 2200 kg m ⁻³
<i>S</i> _z	local longitudinal strain
<i>s</i>	overall strain in a fibre
<i>σ</i>	separation of vibrometer fibre tip and vibrating surface
<i>T</i>	number of twist revolutions imposed on a coupler
<i>t</i>	time
<i>U</i>	waveguide parameter, = $a\sqrt{k^2 n_{core}^2 - \beta^2}$
<i>U</i> ₁ , <i>U</i> ₂	<i>U</i> for each core of a dual-core fibre
<i>u</i>	displacement in flexural acoustic wave
<i>u</i> _o	zero-peak amplitude of flexural acoustic wave
<i>V</i>	waveguide normalised frequency, = $ka\sqrt{n_{core}^2 - n_{cladding}^2}$
<i>V</i> _{co}	<i>V</i> at second-mode cutoff, = 2.405
<i>V</i> _{detected}	detector voltage at the output of the fibre vibrometer
<i>V</i> ₁ , <i>V</i> ₂	<i>V</i> for each core of a dual-core fibre
<i>W</i>	waveguide parameter, = $a\sqrt{\beta^2 - k^2 n_{cladding}^2}$
<i>W</i> ₁ , <i>W</i> ₂	<i>W</i> for each core of a dual-core fibre
<i>X</i>	electrical reactance
<i>X</i> _d	a dual-core fibre parameter
<i>x</i>	taper elongation, transverse co-ordinate perpendicular to plane of symmetry
<i>x</i> ₁ , <i>x</i> ₂	<i>x</i> at the end of the first and second stages of two-part taper elongation
<i>ξ</i>	ratios of powers excited in the modes of a coupler waist, for input in one fibre
<i>y</i>	transverse co-ordinate in the plane of symmetry
<i>Z</i>	electrical impedance, = <i>R</i> + <i>iX</i>
<i>z</i>	longitudinal co-ordinate